



BANK OF ENGLAND

Staff Working Paper No. 650

The effect of house prices on household borrowing: a new approach

James Cloyne, Kilian Huber, Ethan Ilzetzki and Henrik Kleven

February 2017

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Authority Board.



BANK OF ENGLAND

Staff Working Paper No. 650

The effect of house prices on household borrowing: a new approach

James Cloyne,⁽¹⁾ Kilian Huber,⁽²⁾ Ethan Ilzetzki⁽³⁾ and Henrik Kleven⁽⁴⁾

Abstract

We investigate the effect of house prices on household borrowing using administrative mortgage data from the United Kingdom and a new empirical approach. The data contain household-level information on house prices and borrowing in a panel of homeowners, who refinance at regular and quasi-exogenous intervals. The data and setting allow us to develop an empirical approach that exploits house price variation coming from idiosyncratic and exogenous timing of refinance events around the Great Recession. We present two main results. First, there is a clear and robust effect of house prices on borrowing, but the responsiveness is smaller than recent US estimates. Second, the effect of house prices on borrowing can be explained largely by collateral effects. We study the collateral channel in two ways: through a multivariate heterogeneity analysis of proxies for collateral and wealth effects, and through a test that exploits interest rate notches that depend on housing collateral.

Key words: House prices, household borrowing, collateral channel.

JEL classification: D14, E21, E32, E43, E51, G21.

(1) UC Davis. Email: jcloyne@ucdavis.edu

(2) London School of Economics. Email: k.huber@lse.ac.uk

(3) London School of Economics. Email: e.ilzetzki@lse.ac.uk

(4) London School of Economics. Email: h.j.kleven@lse.ac.uk

We thank Orazio Attanasio, Richard Blundell, Christopher Carroll, Eric French, Erzo Luttmer, Magne Mogstad, Emi Nakamura, Sebastian Sieglöcher, David Sraer, Jón Steinsson, and numerous seminar participants for helpful comments and discussions. This research was carried out as part of the Bank of England's One Bank Research Agenda. It uses Financial Conduct Authority (FCA) Product Sales Data that have been provided to the Bank of England under a data-sharing agreement. The FCA Product Sales Data include regulated mortgage contracts only, and therefore exclude other regulated home finance products such as home purchase plans and home reversions, and unregulated products such as second charge lending and buy-to-let mortgages. The views expressed are those of the authors and do not necessarily reflect the views of the Bank of England, the Monetary Policy Committee, the Financial Policy Committee or the Prudential Regulation Authority.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

1 Introduction

It is a well-known fact that house prices are strongly correlated with household borrowing and consumption over the business cycle. These comovements have existed for a long time and were especially strong around the Great Recession. We illustrate this in Figure A.I, which shows the evolution of house price growth, consumption growth, and mortgage debt growth in the US and UK over the last four decades. Motivated by such macro patterns, a leading narrative about the Great Recession argues that house price swings drive borrowing and consumption (e.g. Mian & Sufi 2011, 2014; Mian *et al.* 2013; Kaplan *et al.* 2015; Adelino *et al.* 2016). In this paper we revisit this question using a new approach, providing evidence both on the effect of house prices on borrowing and on the underlying mechanisms driving the effect.

This is an area where causal identification is particularly difficult, because house price variation is endogenous and compelling quasi-experiments are difficult to find. The time series evidence in Figure A.I does not have a causal interpretation, a point emphasized by Campbell & Cocco (2007) and Attanasio *et al.* (2009, 2011). Much of the recent literature instead uses variation in house price growth across geographical areas, which raises concerns about confounding regional shocks (such as shocks to local income expectations) that drive both house prices and the outcome of interest. This requires the use of an instrument for regional house price growth, but fully compelling instruments are difficult to find.¹

Motivated by these challenges, we consider a different setting and a different approach to study the effect of house prices on borrowing. We examine the borrowing decisions of home refinancers using administrative data on the universe of mortgage contracts in the UK from 2005-2015. Our data and setting offer three main advantages. First, the dataset has information on *individual* house prices from mortgage appraisals by lenders. We present evidence showing that, in the UK, mortgage appraisals provide unbiased measures of actual house prices. Second, the data has a *panel dimension* as many homeowners refinance several times during the 10-year window we consider. This results from the fact that refinancing is a frequent phenomenon in the UK, because long-term fixed interest mortgages are not available (see Best *et al.* 2015). The panel dimension of the data

¹Most recent work instruments regional house price growth using a topography-based measure of housing supply elasticities, namely proximity to mountains and oceans that restrict supply (as constructed by Saiz 2010). The idea is that regional housing markets are exposed differently to demand shocks because of their topography. A debate about this instrument highlights potential issues with the exclusion restriction and defiers (see e.g., Davidoff 2013, 2016).

allows us to control for a rich set of fixed effects that deal with the standard confounders discussed in the literature. For example, confounding regional shocks will not be a threat to identification here as we control for county-by-time fixed effects.

Third and finally, the *institutional setting* helps with identification. Most mortgage products in the UK come with a relatively low interest rate for a short time period, typically 2-5 years, followed by a much higher reset rate. This creates a strong incentive to refinance around the onset of the reset rate, and we show that most homeowners do in fact refinance around this time. This implies that the timing of refinance is determined by past contract choices, namely the duration of the initial low interest rate in the last contract.² These mortgage institutions combined with the large house price swings over the period we consider create a potential quasi-experiment. Refinancers face very different house price shocks depending on whether they refinance before, during, or after the housing crisis, and this timing is determined largely by a mortgage contract choice made in the past. Loosely worded, we use the Great Recession interacted with pre-determined, idiosyncratic contract choices as a quasi-experiment for house prices.

We present three main sets of results. The first set of results concerns the impact of house prices on homeowner borrowing. While such borrowing effects are interesting in their own right (e.g. [Mian & Sufi 2011](#)), they are also indicative of the potential consumption effects of house prices and they relate to the same underlying mechanisms as consumption. We find clear evidence that rising house prices induce homeowners to increase borrowing by extracting equity from their home, although the magnitude of the response is smaller than suggested by recent US estimates.³ The elasticity of borrowing with respect to house prices is 0.2-0.3 and is robust across a range of specifications. In our preferred specification, the elasticity is identified from within-individual variation in house price growth. This variation comes from homeowners who refinance at least twice and experience different house price shocks due to how their (quasi-exogenous) refinance timing interacts with the housing cycle.

The second set of results concerns patterns of heterogeneity and mechanisms. The two main reasons why house prices may affect borrowing are wealth effects and collateral effects (see e.g., [Sinai & Souleles 2005](#); [Berger et al. 2015](#)).⁴ All else equal, the wealth effect should be larger for older

²This quasi-exogeneity of refinancing stands in contrast to the US setting where the decision to refinance is endogenous to factors such as income shocks, liquidity needs, and the market interest rate (see [Hurst & Stafford 2004](#)).

³Equity extraction in the UK is similar to home equity loans in the US, with the difference that equity extraction is always rolled into a single home mortgage as opposed to the US institution of secondary mortgages.

⁴A third possible reason is the presence of substitution effects on housing consumption, but this channel is shut down here as we consider refinancers who stay in their existing houses.

homeowners who have short horizons and are therefore in a position to cash in on their housing wealth, while the collateral effect should be larger for more leveraged homeowners. The existing literature has tried to distinguish between different mechanisms by studying such patterns of heterogeneity (Campbell & Cocco 2007; Attanasio *et al.* 2009, 2011). A challenge for such exercises, however, is that different dimensions of heterogeneity are highly correlated. For example, older homeowners have shorter horizons and more asset risk, but are also less levered, and so it is not clear if the age profile is picking up wealth or collateral effects. In this paper we use our rich data to carry out a multivariate and non-parametric analysis of heterogeneity in the elasticity of borrowing with respect to house prices. We consider four dimensions simultaneously — loan-to-value (LTV), borrower age, income, and income growth — and ask how the elasticity varies across bins of a given dimension, holding the other dimensions fixed. The striking finding from this analysis is that there is essentially no heterogeneity in any dimension except one — loan-to-value — but this dimension is very strong. More levered households are much more responsive to house prices, with borrowing elasticities around 0.7 at loan-to-value ratios above 85%. By contrast, the age profile of elasticities is completely flat after controlling non-parametrically for the other dimensions.

The strong relationship between borrowing elasticities and LTV is consistent with evidence on subprime borrowing in the US (Mian & Sufi, 2009, 2011), and it is strongly suggestive of collateral effects. The UK mortgage market offers an interesting way to investigate the collateral channel more deeply, arising from the presence of *observable* credit constraints that depend on collateral. Specifically, as described and analyzed by Best *et al.* (2015), the mortgage interest rate features discrete jumps (notches) at critical LTV thresholds. These notches can be interpreted as soft credit constraints, around which increases in LTV (reductions in collateral) sharply increase the cost of borrowing.⁵ In such a setting it is natural that house price growth leads to larger borrowing by moving homeowners to lower notches and reducing their interest rate, a form of collateral effect.

This brings us to the third set of results, which shows that the borrowing elasticity depends critically on whether the underlying price variation relaxes credit constraints (by pulling homeowners down to lower notches), reinforces credit constraints (by pushing homeowners up to higher notches), or leaves credit constraints unchanged. This analysis shows that the elasticity is much higher (around 0.5) among homeowners whose collateral constraint is relaxed by house price growth, and that the elasticity is zero among those whose collateral constraint is reinforced.

⁵Conceptually, a hard credit constraint (often assumed in models) can be interpreted as the special case of a prohibitively large interest rate notch at a collateral threshold.

Furthermore, we present an analysis of the dynamic interaction between house price growth and bunching responses to interest notches, which is consistent with the collateral channel. This set of findings provide some of the first direct evidence on the collateral channel. Leaving aside questions about identification, most previous work focuses on the total borrowing responses to house prices, because the empirical analyses are based on (regional or aggregate) house price variation that affect both wealth and collateral.⁶

Given that much of the recent literature focuses on the US, it is natural to ask if our results are transportable to the US setting. Two points are worth highlighting. First of all, our empirical design — relying on within-individual variation — identifies micro elasticities rather than macro elasticities. This implies that the various reasons why macro elasticities can vary across economies (such as the underlying source of the house price shock as highlighted by [Kaplan *et al.* 2015](#)) are not relevant for assessing external validity in this case. Second, while micro elasticities may clearly vary across different markets (say because of differences in the strength of credit constraints), we present one piece of evidence in support of external validity: When we use an empirical specification without fixed effects (similar to previous work) and consider the pre-recession period, we obtain borrowing elasticities of a similar magnitude to those found by [Mian & Sufi \(2011\)](#) for the US between 2002-06.

The paper is organized as follows. Section 2 describes the institutional setting and data, section 3 analyzes the sources of house price variation used for identification, section 4 presents results on the effect of house prices on borrowing, section 5 presents results on heterogeneity and mechanisms, and section 6 concludes.

2 Institutional Setting and Data

2.1 UK Mortgage Market

The UK mortgage market has several institutional features that make it an excellent laboratory for investigating the relationship between house prices and homeowner borrowing. In contrast to the US mortgage market, long-term fixed-rate mortgages are unavailable in the UK. The vast majority (more than 90%) of UK mortgage products feature an introductory interest rate followed

⁶An exception is [DeFusco \(2016\)](#) who provides compelling evidence of collateral effects using a natural experiment in Montgomery County, Maryland. Our findings are qualitatively consistent with his, using an entirely different empirical strategy and considering the full population of homeowners with mortgages in the UK.

by a penalizing reset rate.⁷ The initial rate typically has a duration of 2-5 years and this rate may be either fixed or floating. The reset rate lasts for the remainder of the mortgage's duration and is always floating. The reset rate is penalizing in the sense that the same bank almost always offers an identical mortgage product with a lower rate. For example, at current rates a refiner could lower her interest payments by more than 200 basis points (without altering the amortization schedule or other features of the mortgage) by refinancing to avoid the penalizing rate. Besides the penalizing reset after the end of the introductory period, most mortgage contracts feature early repayment charges that make it very costly to refinance or adjust borrowing before the end of the introductory period.

The combination of penalizing reset rates and heavy early repayment charges implies that households have strong incentives to refinance at the end of the initial duration. To confirm that households act on these incentives, Figure 1 shows the distribution of time between mortgages among refiners in our data. The distribution features large spikes in refinancing activity around 2, 3, and 5 years after the previous mortgage, consistent with the fact that these are the most common durations on offer. The lightly shaded bars indicate the fraction of households in each month that refinance around the end date of their initial low-interest duration (within a window of 2 months before and 6 months after the end date). The figure demonstrates that the vast majority of households refinance around the time that the initial duration ends.⁸

This institutional setting has the following key advantages for our empirical approach. First, the fact that refinancing occurs around predetermined dates makes the time of refinance potentially orthogonal to individual circumstances. This contrasts with the US setting where the decision to refinance or take out home equity loans is likely to be correlated with unusual consumption and borrowing needs (see [Hurst & Stafford 2004](#)). Second, the fact that refinance events are frequent allows us to observe the same homeowner refinancing several times, facilitating the use of panel data methods. Third, the frequency of refinancing also implies that the market for home equity loans is minimal in the UK. As households are only a few years away from refinancing at any given time, home-equity based borrowing is done almost exclusively through equity extraction at the time of refinancing. Finally, it is worth highlighting that mortgage debt comprises nearly 90%

⁷The prevalence of products with an introductory rate can be seen by looking at publications from MoneyFacts.co.uk, a company that collects and publishes information on all mortgage products on the market each month.

⁸How do borrowers choose their mortgage's initial duration? The main determinants in this choice are interest rates and expectations thereof. For example, a two-year initial duration will offer a lower interest rate than a five-year initial duration, but the five-year product hedges against interest rate increases in the remaining three years. The choice between the two will be determined by, among other things, risk preferences. Our empirical approach will be able to deal with unobserved heterogeneity in preferences for low-interest durations.

of all household debt in the UK. Thus studying borrowing responses in the mortgage market gives a nearly complete view of household borrowing behavior.

When households refinance, the lender appraises the house value and this appraisal determines home equity. The household's decision about equity extraction then determines the new debt level, the loan-to-value (LTV) ratio, and the interest rate. The interest rate charged on UK mortgages follows a step function with discrete jumps (notches) at certain LTV thresholds. The most common interest rate notches occur at LTVs of 60%, 70%, 75%, 80%, and 85%. Figure A.II in the appendix shows the average interest rate schedule as a function of LTV across all mortgage products (see [Best et al. 2015](#) for details).⁹ The overall level of the interest rate schedule depends on a number of mortgage contract characteristics (including the duration of the initial interest rate), but all contracts feature notches at critical LTV thresholds. These interest notches introduce a form of 'soft' credit constraints that depend on collateral values: borrowing costs jump sharply as the LTV ratio exceeds — and the collateral therefore falls below — the critical thresholds.¹⁰ House price growth reduces a homeowners's LTV ratio, allowing her to borrow at a lower interest rate if it pulls her across interest notches. We will utilize this institutional feature to devise a test for the collateral channel.

2.2 House Price Measurement

We measure house prices based on lenders' house value appraisals. There are a number of useful reasons for this. First, these appraisals provide us with house price information at the individual level. Second, appraisals take place at every refinance event, providing us with several observations of house prices for each house-homeowner pair. Third, the appraisal provides the exact house price measure used by the lender to determine collateral, the LTV ratio and the interest rate. Hence, for capturing the collateral effect of house prices, there is no measurement error in the price measure we use.

Nevertheless, a potential concern with our house price measure is the presence of appraisal bias. A literature has shown that mortgage appraisals feature systematic upward bias in the US (e.g. [Ben-David 2011](#); [Agarwal et al. 2015, 2016](#)), which may reduce the suitability of appraisals for capturing the true wealth effect of house prices in that setting. However, such appraisal bias does not seem to be a problem in the UK, as we demonstrate in two ways. First, while we do not

⁹[Best et al. \(2015\)](#) provide a bunching analysis of borrowing responses to these interest rate notches.

¹⁰Alongside these notches, there is also a hard collateral constraint as only a handful of mortgage products are currently available at LTVs exceeding 90%.

observe actual market prices for refinanced properties, we do observe market prices (along with appraisals) when properties are purchased and the first mortgage is originated. Hence, Figure 2 shows a histogram of the difference between the purchase price and the appraisal for transacted properties. The difference is zero for the vast majority of transactions, showing that appraisals line up with the actual price for newly purchased homes.

However, appraisal bias may be more acute for refinances than for first mortgages, as there is no purchase price to anchor the appraisal for refinances. This motivates our second test in which we compare actual purchase prices (for transacted properties) with appraised prices (for refinanced properties) over time. The results are shown in Figure 3. Panel A plots the raw time series of actual and appraised prices. Taken at face value, this panel suggests that there *is* bias: appraised prices are slightly higher than purchase prices on average and the appraisals appear too smooth during the financial crisis. But such a comparison does not account for the fact that the composition of properties in the two series is different, and that the composition of each series changes over time. To be able to accurately compare the two series and their changes over time, Panel B presents regression-adjusted price series in which we control non-parametrically for two observables: the age of the homeowner and the postcode of the property. Specifically, we run the following regression separately for the price and the appraisal series:

$$P_i = \sum_t \beta_t \cdot \mathbf{I}[quarter_i \in t] + \sum_k \gamma_k \cdot \mathbf{I}[age_i \in k] + \sum_p \lambda_p \cdot \mathbf{I}[postcode_i \in p] + \nu_i, \quad (1)$$

where the first term includes a full set of quarter dummies, the second term includes dummies for twenty quantiles of the age distribution, and the third term includes dummies for twenty quantiles of the postcode-level distribution of house prices. Specifically, the last term is based on the average house price of each 6-digit postcode, including dummies for the postcode's quantile position in the distribution of postcode-level prices. This term controls for the fact that the quality of neighborhoods that feature high or low activity differs across the two series and changes over time.

The plotted values in Panel B are the coefficients on the quarter dummies from equation (1), adding a constant equal to the effect of the average age and the average postcode (in each series separately). We see that, with non-parametric controls only for age and neighborhood, the two series track each other closely throughout the period and the recession is now clearly visible in the appraisal series. In other words, the differences in Panel A were due to differences in sample

composition rather than real appraisal bias. We therefore conclude that appraisals are a good reflection of true property prices in the UK market.

2.3 Data

The data come from a novel and comprehensive regulatory dataset containing the universe of mortgage product sales. These data are collected by the UK's Financial Conduct Authority (FCA) and available to restricted members of staff at the FCA and the Bank of England. This Product Sales Database (PSD) has information on all completed household mortgage product originations from April 2005, but does not include commercial or buy-to-let mortgages.¹¹

Regulated lenders are required to submit quarterly information on all mortgage originations. The data include a range of information about the mortgage such as the loan size, the date the mortgage became active, the house price appraisal, the interest rate charged during the introductory period, whether the interest rate is fixed or variable, the end date of the initial duration (the time at which the higher reset rate starts applying), whether mortgage payments include amortization, and the mortgage term over which the full loan will be repaid. The data also include a number of borrower characteristics such as borrower age, gross income, and whether the income is solely or jointly earned.¹²

Another useful feature of the PSD is that it contains information on whether the household is a refiner. Using information about the characteristics of the property and the borrower, refinancing households can be matched over time to construct a panel. As noted above, since refinancing is a regular occurrence in the UK mortgage market, this provides us with multiple observations for the same household over the 11 years of the sample. Using our new panel, we can compute a range of useful household-level statistics including house price growth, mortgage debt growth, amortization, and equity extraction/injection.

Overall, the PSD contains around 14 million mortgage observations. Around half of these observations are mortgages for new house purchases, while the other half are refinancing events. Since we need to calculate the house price change and equity extracted for our analysis, we can only use refinancing observations where we observe a previous mortgage event (either the house purchase or a previous refinancing event) by the same household for the same property. Our sample is therefore a subset of the refiners in the PSD.

¹¹See <https://www.fca.org.uk/firms/product-sales-data> for officially published high level data.

¹²Full details of the dataset can be found on the FCA's PSD website.

Table 1 summarizes the data. Panel A compares descriptive statistics for home buyers (column 1), all refinancers (column 2), and refinancers in our estimation sample (column 3). There are no significant differences between the three groups in the share of couples, income, income growth, interest rate, and house price. Some differences between buyers and refinancers are to be expected. For example, buyers tend to be younger and have higher LTV ratios.

Panel B of Table 1 reports statistics for the 1.38 million observations in our estimation sample, split into three subsamples. As discussed above, practically all mortgages in the UK have an initial duration with a favorable interest rate, after which a higher reset rate kicks in. This gives a strong incentive for refinancing around the onset of the reset rate. The subsample in column 1 of panel B includes the 0.48 million observations where we know refinancing took place “on-time” (defined as between 2 months before and 6 months after the reset rate onset), while column 2 includes the 0.28 million observations where we know refinancing took place “off-time”. For a large part of the sample, 0.61 million observations, we do not observe when the reset rate kicks in, because lenders were not always required to report this statistic to the Financial Conduct Authority. We summarize these observations in column 3. There are no significant differences across the three groups in any of the observables.

3 House Price Variation

There is large house price variation in the data. Figure 4 shows the distribution of house price growth between refinance events for homeowners in our estimation sample. To measure individual house price growth, the sample conditions on observing homeowners at least twice. The first price observation for each homeowner may come either from the first mortgage in the house or a refinance, while subsequent price observations always come from refinances. The distribution shows that house price growth lies between -30% and +60% across refinance events, giving us lots of variation to work with. We note that there is some round-number bunching at zero price growth, suggesting that some lenders set the new house price equal to the old house price whenever the two are very close (see [Kleven 2016](#) for a discussion of round-number bunching).

While there is large house price variation in the data, the challenge is that much of it may be endogenous to demand factors that impact our outcome of interest. Our approach starts by controlling for obvious confounders by absorbing a rich set of fixed effects. Individual fixed effects control for time-invariant individual preferences for borrowing, month fixed effects control

for time-varying macro factors that affect borrowing, while county-by-year fixed effects control for local, time-varying shocks to borrowing demand. Specifically, ‘counties’ are defined as local planning authorities (or councils), of which there are more than 400 in the UK and 32 in London alone.

Figure 5 shows the distribution of residual house price growth, after absorbing the fixed effects described above. Allowing for individual fixed effects on house price growth gives an R-squared of one among households with just two mortgage observations (one price growth observation), so the figure considers the sample of homeowners observed at least three times. Panel A shows the raw distribution of house price growth in this subsample as a benchmark (it looks similar to the raw distribution in the previous figure), while Panel B shows the residualized distribution. Importantly, there is large remaining house price variation even after controlling for fixed effects, between -20% and +20% across refinance events.

What drives this residual variation? In general there can be two sources of remaining variation. The first is that different properties experience different price growth *within counties*, so that county-by-year fixed effects do not fully absorb the housing cycle. This arises because of variation across neighborhoods within counties, variation across property types within neighborhoods, or completely idiosyncratic variation driven by features of the specific house. On the latter, note that the value of a specific house may increase due to home improvements undertaken by the owner, which would not be real house price appreciation. However, the data include an indicator for home improvement activity, which allow us to deal with this potential issue. Moreover, as described below, we consider IV-specifications that are unlikely to be affected by home improvements.

The second source of variation is idiosyncratic variation in the timing of refinance events relative to the price cycle. As described above, homeowners have a strong incentive to refinance around the onset of the reset rate, typically after 2, 3 or 5 years, as these are the most common products in the market. Hence, the timing of refinance is determined to a large extent by a duration choice made several years in advance, creating arguably quasi-exogenous variation. Figure 6 illustrates conceptually how this works. It compares two homeowners who start out at the same time (time 0), live in houses with the same price cycle (the solid blue line), but have different preferences over low-interest rate durations. One homeowner prefers 2-year fixed interest rate loans, while the other prefers 3-year fixed interest loans. Of course, this difference in duration preferences will be related to, for example, risk preferences that may themselves impact on borrowing

behavior, but such time-invariant preference heterogeneity is absorbed by the individual fixed effect. What creates variation here is the *interaction* of idiosyncratic duration preferences with the housing cycle: The 2-year person refinances three times over a 6-year period, facing either positive or negative price shocks at each event, whereas the 3-year person refinances only two times facing a zero price shock each time. Our empirical strategy exploits this kind of within-person variation in price growth.

In Figure 7 we illustrate this point using the actual data. The figure plots average house price growth for homeowners who refinance at different times (in January of different years) by bins of the duration of their last mortgage. The two panels show the same graphs, but highlight two different homeowners who experience very different within-person price patterns due to past duration choices. The homeowner in Panel A refinances in January 2010 coming out of a 2-year mortgage chosen in 2008, and refinances again in January 2013 coming out of a 3-year mortgage chosen in 2010. This homeowner experiences a substantial negative shock the first time around, and a substantial positive shock the second time around. The homeowner in Panel B also refinances in January 2010 and January 2013, with the only difference being that in 2010 she was coming out of a 5-year mortgage chosen in 2005. As a result, this homeowner faces similar positive price growth in both refinance events. The empirical approach we propose uses this kind of within-person variation for identification: i.e., we use the change over time for Person A (who goes from negative to positive price growth) relative to the change over time for Person B (who goes from positive to positive price growth). This is a form of *triple-differences* strategy as we are comparing within-person changes in price *growth*.

The exogeneity of this duration-driven variation in house price growth requires that homeowners are not choosing durations in anticipation of future house price growth and future borrowing needs. For example, if homeowners were choosing 2-year mortgages (rather than 3-year mortgages) in late 2005 — anticipating that this would put them at the peak of the boom (rather than at the bottom of the bust) — to be able to extract more equity for consumption goods in late 2007, then our estimates would not be causally identified. A sufficient condition for ruling out such hyper-rational and forward-looking behavior is that homeowners are not able to forecast house prices with much precision. This assumption seems particularly persuasive around the time of the Great Recession, and it is consistent with a growing consensus that homeowners tend to have biased beliefs about future house prices (e.g., [Case & Shiller 1989](#); [Shiller 2007](#); [Case et al. 2012](#); [Kaplan et al. 2015](#)). However, we do not necessarily need bias or irrationality for our strategy to

work; a sufficient amount of house price uncertainty will do.

Another way of gauging the exogeneity of duration-driven house price growth is to check if duration choices, besides predicting future house price appreciation, predict other things of relevance to borrowing. Hence, Figure A.III in the appendix shows how much of the residual price variation (Panel A) and residual income variation (Panel B) can be explained by past duration choices, having absorbed all the other fixed effects. The figure shows that, while duration choices are strong predictors of future price growth, they do not predict future income. This lends further support to our strategy.

We present results from two types of strategies. We first consider OLS fixed effects regressions, which use all of the residual variation for identification. This includes idiosyncratic variation in price growth across properties within counties, and it includes idiosyncratic variation in the timing of refinance events. As discussed earlier, a concern with the first source of variation is that it may be partly driven by home improvements. Hence, we also consider IV-regressions in which we construct instruments based on past duration choices (which determine refinance timing). These results should not be affected by home improvements. Reassuringly, our OLS fixed effects and IV results turn out to be quite similar.

4 Do House Prices Affect Borrowing?

4.1 Baseline Specification

To establish a baseline, we start from a specification that is similar in spirit to specifications used in existing work. Specifically, we consider the following specification

$$\Delta \log D_{it} = \sum_j \beta_j \cdot \mathbf{I}[\Delta \log P_{it} \in j] + \nu_{it}, \quad (2)$$

where D_{it} and P_{it} denote mortgage debt and house prices, respectively, for individual i at time t . While we primarily consider log-specifications, we will also explore level-specifications and show that those yield the same qualitative results.¹³ We allow for different bins of house price growth to have different effects on borrowing, as we do not (yet) want to commit to a specific functional form.

¹³To be clear, the coefficient obtained from a log-specification represents a borrowing elasticity, whereas the coefficient obtained from a level-specification represents a marginal propensity to borrow (which can be used to assess the average borrowing elasticity in order to compare with the log-specification).

Equation (2) corresponds to the specification used by [Mian & Sufi \(2011\)](#), except for three differences: (i) we rely on individual rather than regional house price variation, (ii) we allow for a non-parametric specification without a priori functional form restrictions, and (iii) we do not instrument house prices using topography-based housing supply elasticities.¹⁴

Panel A of Figure 8 shows the results from this specification. It plots the log-change in mortgage debt against bins of the log-change in house prices (where the log-change is scaled to approximate the percentage change). Three insights are worth highlighting. First, overall there is a clear positive relationship between house price growth and debt growth. Debt growth changes from 0% to 15% as house price growth changes from -10% to +40%. Second, there is a strong asymmetry between negative and positive price shocks: Homeowners increase debt when their house becomes more valuable, but they do not reduce debt when their house becomes less valuable. A possible explanation for this phenomenon is the presence of liquidity constraints that prevent homeowners from injecting equity when negative house price shocks push up their LTV ratios. Third, the average elasticity of borrowing across the full range of house price growth — obtained from a log-linear specification — equals 0.3. However, this elasticity masks the heterogeneity between the negative and positive ranges, with an elasticity of 0 in negative ranges and an elasticity of 0.4 in positive ranges.

Panel B of Figure 8 investigates cyclical variation in the elasticity of borrowing. Again, we consider the elasticity obtained from a log-linear specification, splitting the sample into different years. The graph shows that the borrowing elasticity is strongly pro-cyclical, with the largest elasticities in the run-up to the recession and the smallest elasticities in the middle of the recession. This elasticity cycle is consistent with the asymmetry between negative and positive shocks shown in Panel A, but it turns out that this is far from the whole story: When we condition the sample on positive price growth, or add a fixed effect for negative price growth in the full sample, the cyclical elasticity largely survives. In general, there can be two possible explanations for the elasticity cycle. The first possibility is that the true elasticity is cyclical, because the underlying mechanisms driving the effect change over the business cycle. The second possibility is that the true elasticity is not cyclical, but rather that a cyclical omitted variable (e.g. expectations) is creating a spurious cyclical estimate. As we move to better identified specifications below, we will be able to distinguish

¹⁴[Hilber & Vermeulen \(2016\)](#) constructs a topography-based housing supply elasticity index for England (a la [Saiz 2010](#)), but not for the rest of the UK (Northern Ireland, Scotland and Wales). However, besides the potential issues with the exclusion restriction of such instruments (as discussed in the introduction), [Hilber & Vermeulen \(2016\)](#) show that the instrument does not have a strong first stage in the English setting: Topography does not predict house price variation in this country.

between these two hypotheses.

The outcome considered above equals total debt growth between the current and the last refinance event, i.e. $\Delta \log D_{it} = \log D_{it} - \log D_{it-1}$. This outcome captures both the equity extraction decision made by the homeowner at time t and the amortization between times $t - 1$ and t . Because the amortization schedule was chosen as part of the last mortgage contract, it cannot respond to house price appreciation between $t - 1$ and t . If there is any spurious correlation between amortization schedules and house price variation, this will lead to bias. Fortunately, we have sufficiently detailed information about mortgage contracts to precisely assess amortization for each homeowner (see also [Best et al. 2015](#)). Hence, we now turn to an improved outcome variable that captures only the active equity extraction decision made at time t . This outcome equals $\log D_{it} - \log D_{it}^P$, where D_{it}^P denotes the pre-determined debt at time t based on past debt choices and amortization.¹⁵

The results for this outcome are shown in Figure 9, which is constructed exactly as is the previous figure. The results are qualitatively similar: there is a clear positive relationship between house price growth and borrowing overall, there is a strong asymmetry between positive and negative house price growth, and the borrowing elasticity is cyclical. The average elasticity of borrowing is a bit smaller here (0.23), but this reflects that the slope is now negative (as opposed to zero before) in the range of falling house prices. Within the range of increasing house prices, the elasticity is about 0.4, corresponding to the finding in the previous figure. In the analysis that follows we consider equity extraction as our outcome.

4.2 Fixed Effects Specification

Taking advantage of the fact that we observe multiple refinance events for each individual, we augment the baseline specification with fixed effects. That is, we specify

$$\log D_{ict} - \log D_{ict}^P = \sum_j \beta_j \cdot \mathbf{I}[\Delta \log P_{ict} \in j] + \alpha_i + \gamma_t + \delta_{ct} + \nu_{ict}, \quad (3)$$

where the index c denotes county (local planning authority as described above), α_i is the individual fixed effect, γ_t is a time fixed effect (at the monthly level), and δ_{ct} is a county-by-time fixed effect (at the yearly level).¹⁶ The county-by-time fixed effect absorbs regional, time-varying factors (such

¹⁵That is, we have $D_{it}^P = D_{it-1} + \text{amortization between } t - 1 \text{ and } t$.

¹⁶There is obviously some abuse of notation here as we use t to describe time in both months and years.

as local shocks to income expectations), thus dealing directly with the main confounder discussed in the previous literature. By allowing for individual fixed effects in a first-differenced equation, this has the form of a triple-differences specification relying on within-individual variation in price growth.

The results are shown in Figure 10, distinguishing between a specification with individual and time fixed effects only (Panel A) and a specification that adds county-by-time fixed effects (Panel B). The following insights emerge. First, the two different fixed-effects specifications yield almost identical results. Once we have controlled for individual and time fixed effects, adding county-time fixed effects have no noticeable effect. Second, the relationship between equity extraction and house price growth is now monotonically increasing and almost perfectly linear in logs. There is no longer any asymmetry between negative and positive shocks. Third, the borrowing elasticity is now smaller, about 0.2.

The results in Figure 10 could be biased by individual, time-varying effects that are correlated with house price growth and impact on borrowing behavior. To investigate this threat to identification, we can exploit that the data include information on a number of individual, time-varying variables that are relevant for debt demand. Hence we consider a specification with such individual controls:

$$\log D_{ict} - \log D_{ict}^P = \sum_j \beta_j \cdot \mathbf{I}[\Delta \log P_{ict} \in j] + \alpha_i + \gamma_t + \delta_{ct} + X_{it}\theta + \nu_{ict}, \quad (4)$$

where X_{it} includes the income level, the income growth, the last mortgage interest rate, the age of the borrower, a dummy for couples, and dummies for a range of self-reported reasons for the current and the last refinance (including home improvement as a key category).

The results are shown in Figure 11. Panel A of the figure shows that the results are completely stable when moving to this richer specification: the average elasticity is the same and in fact the entire functional form is the same. Panel B of the figure returns to the question of cyclicity in the elasticity. The elasticities reported in this figure are based on a log-linear version of equation (4), interacting the house price growth variable by year dummies. We see that the richer specification has eliminated some, but not all of the elasticity cycle shown earlier. Hence we conclude that the strong cycle observed for the baseline specification was partly a result of bias and partly a real phenomenon.

As discussed in sections 2 and 3, our empirical strategy is based on the idea that the timing

of refinance is quasi-exogenous in the UK. The argument was that homeowners tend to refinance around the onset of the reset rate, the timing of which is determined by a duration choice made in the last refinance event. We showed in section 2.1 that a majority of homeowners do indeed refinance around the onset of the reset rate, but we also saw that some homeowners refinance at other times, typically ‘too late’. There are a variety of reasons why some homeowners might refinance late — including inattention and financial distress — but whatever the reason, it raises the concern that such homeowners endogenously tailor the timing of refinance to house price movements. If this is so, our estimates based on the full sample of refinancers — including both on-time and off-time refinancers — may be subject to selection bias.

To investigate this selection issue, Table 2 presents estimates of borrowing elasticities across samples that vary by refinance timing: the full sample in Panel A (summarizing the results already presented), the sample of on-time refinancers in Panel B, the sample of off-time refinancers in Panel C, and the sample of refinancers with missing duration information in Panel D. As mentioned earlier, even though almost all mortgage contracts in the UK (including variable rate loans) come with a penalizing reset rate after a certain duration, we do not observe this duration for all homeowners as it was not always mandatory for lenders to provide it.¹⁷ Overall, the table shows that elasticity estimates are very robust: across all four samples and fixed-effects specifications (columns 2-4), the elasticity varies between 0.17 and 0.27. It is interesting, however, that the elasticity is consistently higher in the off-time sample, consistent with a small selection bias. Hence our preferred estimates are those based on the on-time sample, featuring borrowing elasticities that are slightly smaller than those reported above.

Finally, we present two additional robustness checks on the fixed effects specification. The first check investigates the issue of home improvements. While working with individual house price information has several advantages, they do introduce a problem not present in regional-level data: the house price variation may be driven partly by idiosyncratic home improvements, which do not represent true increases in household net worth.¹⁸ As a first check we use the fact that the data contain information on the reason for refinancing, including home improvements as one category. Hence, for homeowners who reported home improvement in their last refinance, we

¹⁷To be clear, we always observe the *actual* time between refinance events, it is only the duration of the low-interest rate period defined in the mortgage contract that we do not always observe. In the sample of homeowners with missing duration information, the actual time between refinance events features strong bunching at 2, 3 and 5 years, showing that these households do in fact have a fixed low-interest duration.

¹⁸That is, home improvements do not increase household net worth unless the market value of the house increases by more than the amount invested in the house. Besides this, there are of course endogeneity concerns with home improvement-related house price appreciation.

know that house price growth in the current refinance is likely to be driven partly by home improvements. Hence, Table A.I shows elasticity estimates in three subsamples: homeowners whose last refinance was for home improvements (Panel A), homeowners whose last refinance was not for home improvements (Panel B), and homeowners for whom the reason for the last refinance is unknown. The table shows that, for all the fixed effects specifications, the estimated elasticity is quite stable across samples. Specifically, among those who report no home improvement, we find elasticities that are similar to the elasticities for the full sample discussed above. While this alleviates any major concerns about home improvements, the next section goes further by presenting IV-estimates that cannot be plausibly affected by home improvements.

The second check investigates alternative functional forms. Starting from the specification in equation (3), Figure A.IV shows how the results are affected by moving from a log-specification to a level-specification (Panel A) and by moving in addition from house prices to housing net worth as the explanatory variable (Panel B). In each panel we continue to allow for different bins of the explanatory variable to have different effects on borrowing. While the results are qualitatively unaffected by these changes, the alternative specifications are useful for obtaining different types of parameters. Panel A yields an estimate of the marginal propensity to borrow (equal to 0.11) as opposed to the elasticity parameter discussed so far. Panel B yields an estimate of the elasticity with respect to housing net worth — defined as house price minus baseline mortgage debt — as opposed to the elasticity with respect to house prices.¹⁹ The fact that the elasticity with respect to housing net worth is considerably smaller is a mechanical rather than substantive result: because housing net worth is only a fraction of the house price, any given log-change in house prices translates into a much larger log-change in housing net worth. This makes the elasticity with respect to housing net worth mechanically smaller.

4.3 IV Specification

Our fixed effects specification relies on two sources of residual variation: (i) idiosyncratic variation in price growth across houses within counties, (ii) idiosyncratic variation in the timing of refinance events across homeowners. The first source of variation could be endogenous due to for example home improvements (as discussed above) or endogenous selection into neighborhoods. Hence, in this section we consider an IV-strategy that relies solely on variation in the (pre-determined)

¹⁹We specify housing net worth as the house price minus *baseline* debt (as opposed to current debt) in order to avoid a clear endogeneity problem. This implies that the variation in housing net worth comes from the variation in house prices, and so the two elasticities are identified from the same source of variation.

timing of refinance events.

We do not want to rely on cross-sectional variation in duration choices, because these are insurance choices that reflect risk preferences and therefore may affect borrowing directly. As discussed above (see Figure 7), the most compelling source of variation is the *interaction* between the duration choice in the last mortgage (say 2-year vs 3-year fixed interest rate) and the time of the current refinance event (say 2010 vs 2011). Hence we construct instruments based on the interaction between dummies for past duration choices and dummies for the time of refinance, within different regions. The first stage of the IV is specified as follows

$$\Delta \log P_{ict} = \rho \cdot \text{last duration}_{it} \otimes \text{year}_t \otimes \text{region}_i + \alpha_i + \gamma_t + \delta_{ct} + X_{it}\eta + \mu_{ict}, \quad (5)$$

where \otimes denotes the outer product, so that the instrumental variables ($\text{last duration}_{it} \otimes \text{year}_t \otimes \text{region}_i$) include every possible interaction between last duration dummies, year of refinance dummies, and regional dummies. It is for computational reasons that the instruments are based on year dummies (rather than month dummies) and region dummies (rather than the more disaggregated county dummies). We include fixed effects for household, month and county-by-year, and we also allow for individual, time-varying controls X_{it} including duration dummies on their own. This specification implies we are identifying off of the interaction between last duration and time, taking out the average effects of duration and time separately.

The second stage of the IV is similar to the fixed effect specifications considered earlier, i.e.

$$\log D_{ict} - \log D_{ict}^P = \beta \cdot \Delta \widehat{\log P}_{ict} + \alpha_i + \gamma_t + \delta_{ct} + X_{it}\theta + \nu_{ict}, \quad (6)$$

where $\Delta \widehat{\log P}_{ict}$ is the predicted house price growth from the first-stage specification (5).

The results are shown in Table 3. The table shows the estimated elasticities of equity extraction with respect to house price across five IV specifications. The richest specification in column (5) corresponds to the specification shown in equations (5)-(6). There is a non-trivial difference in the estimates between the basic specifications without household and month fixed effects (columns (1)-(2)) and the richer specification with those fixed effects (columns (3)-(5)). But across the richer specifications, the IV elasticity estimates are very stable (around 0.28-0.29) and slightly higher than the OLS estimates shown earlier. The fact that the IV estimates are higher is consistent with a (small) bias from home improvements in the OLS estimates: house price appreciation due to home improvements does not represent real appreciation and would therefore tend to attenuate the OLS

estimates. These differences notwithstanding, the IV table confirms the overall qualitative results presented so far: There is a clear positive effect of house prices on borrowing, but the effects are smaller than recent estimates have suggested.

5 Why Do House Prices Affect Borrowing?

Having established a causal relationship between house prices and household borrowing, we now investigate the reasons for this relationship. [Berger *et al.* \(2015\)](#) provide a theoretical foundation for the various mechanisms that may be at play. Here we focus on the two main mechanisms discussed in the literature.

First, higher house prices increase homeowners' nominal housing wealth, so that borrowing responses may reflect the marginal propensity to consume out of wealth ([Campbell & Cocco 2007](#); [Case *et al.* 2013](#)). However, it is not obvious that such changes in nominal wealth translate into real wealth, as highlighted by [Sinai & Souleles \(2005\)](#). They argue that homeownership provides a hedge against future housing expenditures for households with long expected tenures in their existing homes. This implies that house prices have negligible effects on lifetime net worth and should not affect borrowing. If wealth effects are operational they must therefore rely on expected changes in real housing consumption over the lifecycle. For example, old homeowners may expect to downsize or exit the housing market in the near future, in which case house price growth tends to increase net wealth. Young homeowners, on the other hand, have constant or increasing housing needs over the foreseeable future, so that the nominal wealth effect of house price growth will be offset by increases in future housing expenditures. This suggests larger wealth effects for old homeowners than for young homeowners. Hence, a number of existing papers assess the importance of wealth effects by studying heterogeneity with respect to age, but with conflicting results ([Attanasio & Weber 1994](#); [Campbell & Cocco 2007](#); [Attanasio *et al.* 2009](#); [Mian & Sufi 2011](#)).

Second, housing wealth is the largest form of household collateral. An increase in nominal housing wealth may therefore relax borrowing constraints, which tend to be proportional to collateral values. The collateral channel has been studied theoretically in the macro housing literature (e.g., [Aoki *et al.* 2004](#); [Iacoviello 2005](#)), and it has been argued to be empirically important for household borrowing in a number of studies (e.g., [Lustig & Nieuwerburgh 2005](#); [Mian & Sufi 2011](#); [DeFusco 2016](#)).²⁰ The collateral channel implies heterogeneity across leverage ratios: Households

²⁰The collateral channel has also been shown to be important for business investments and employment ([Chaney *et al.* 2012](#); [Adelino *et al.* 2015](#)).

with higher leverage are more collateral constrained, and house price appreciation is therefore more likely to relax collateral constraints for such households.

In the next section, we take a first step towards disentangling wealth and collateral effects based on a heterogeneity analysis that uses the power and granularity of our administrative data. We analyze heterogeneity in the borrowing elasticity along the main dimensions predicted to determine household borrowing responses, including borrower age and leverage. This analysis suggests that the collateral channel plays a crucial role. We then explore the collateral channel more closely in the following section, proposing a new method to assess its empirical importance.

5.1 Heterogeneity Analysis

We investigate how borrowing elasticities vary along four dimensions of heterogeneity: loan-to-value (LTV), borrower age, income level, and income growth. We consider two types of specifications. Univariate specifications investigate heterogeneity in each dimension separately, while multivariate specifications allow for heterogeneity in all four dimensions simultaneously. Many dimensions of heterogeneity are highly correlated, making it difficult to interpret results from univariate heterogeneity analyses. Our multivariate specifications allow us to disentangle which dimensions truly drive heterogeneity in responsiveness, and which dimensions only appear to do so by being correlated with other relevant dimensions. We estimate specifications of the type

$$\log D_{it} - \log D_{it}^P = \sum_k \sum_j \beta_j^k \cdot \mathbf{I} [X_{it}^k \in j] \cdot \Delta \log P_{it} + \sum_k \sum_j \lambda_j^k \cdot \mathbf{I} [X_{it}^k \in j] + \nu_{it}, \quad (7)$$

where $\mathbf{I} [X_{it}^k \in j]$ is a dummy equal to one when variable k (LTV, borrower age, income, or income growth) falls in bin j . By allowing for a large set of bin dummies in each dimension (7 LTV bins, 9 age bins, 7 income bins, and 7 income growth bins), and by allowing for these dummies to affect both the slope and the intercept, our analysis is very non-parametric. Hence, the heterogeneity patterns we uncover will not be driven by overly restrictive functional form assumptions. We do assume that the effect of prices on borrowing *within* dimension k and bin j is log-linear, but this assumption is a good approximation as we show below. To increase precision, specification (7) does not include the household and time fixed effects considered in the previous section. It is possible to consider such an extension and the heterogeneity results turn out to be very similar, but standard errors increase substantially in fixed effects specifications with heterogeneity.

In Figure 12 we investigate heterogeneity with respect to age and LTV, which are the two main

proxies for wealth and collateral effects as discussed above. The top panels show heterogeneity by pre-determined LTV, defined as the LTV ratio absent any equity extraction/injection and absent any house price growth between the current and last refinance event. This LTV is determined by the last choice of mortgage debt and amortization along with the last house price. The graphs show a strong monotonic relationship between the borrowing elasticity and LTV. This holds both when studying this dimension of heterogeneity on its own (Panel A) and when controlling for the other dimensions of heterogeneity (Panel B). In fact, going from the univariate to the multivariate specification hardly affects the relationship, although it increases standard errors somewhat. Hence, homeowners with low levels of collateral borrow much more against house price increases than do those with high levels of collateral. The strong degree of LTV heterogeneity is not driven by the log-linearity assumption made in equation (7), which we show in a fully non-parametric specification in appendix Figure A.V.²¹

The bottom panels of Figure 12 investigate the effect of age. These panels show heterogeneity in the borrowing elasticity across 5-year bins between the ages of 20 and 60. Panel C presents results without controls for the other dimensions of heterogeneity. The figure shows the opposite pattern from the wealth effects suggested by standard lifecycle theory: young households are more responsive to house prices than old households. A similar pattern of heterogeneity was found by *Attanasio et al. (2009)* using UK survey data and structural methods. They suggest that this puzzling pattern might arise because the young tend to be more leveraged than the old, so that collateral effects confound wealth effects (see *Berger et al. 2015* for a similar argument). Panel D investigates and confirms this hypothesis. It shows that, once we control for LTV (as well as income and income growth), the age profile of borrowing elasticities is completely flat.

For completeness, Figure 13 displays heterogeneity across income levels (top panels) and income growth (bottom panels). Income is measured at the time of the last refinancing event, while income growth is measured as the log-change since the last refinancing event. We use dummies representing seven quantiles of the distribution of each of these variables. Once again, we consider the univariate specification on the left and the multivariate specification on the right. These graphs

²¹Figure A.V presents non-parametric estimates allowing for a large set of bin dummies for house price growth (as in the previous section) within three separate LTV categories. The three samples correspond to low-leverage homeowners (LTV below 60%), intermediate-leverage homeowners (LTV between 60-80%), and high-leverage homeowners (LTV above 80%). Two insights are worth highlighting. First, the *level* of equity extraction decreases with leverage as one might expect: highly leveraged households have a larger stock of existing debt, are more constrained in their borrowing capacity, and should be on an amortization path over their lifecycle. Second, the *slope* of equity extraction increases with leverage, consistent with our previous findings on elasticity heterogeneity. That is, homeowners with high leverage (low collateral) extract less equity, but are more inclined to increase equity extraction when house prices go up.

do not show any noticeable patterns of heterogeneity: they are quite flat across both income levels and income growth in both the univariate and multivariate cases.

How should we interpret these heterogeneity patterns? The fact that leverage is such a strong predictor of borrowing elasticities, even after controlling non-parametrically for other correlated factors, while at the same time the other factors have no predictive power clearly points to the collateral channel as being central. A few qualifications to this interpretation are worth mentioning. First, wealth effects may not be the only force driving heterogeneity across borrower age (even conditional on the other controls), and so the flat age profile does not rule out wealth effects. Second, wealth effects may themselves lead to heterogeneity across LTV ratios, even absent a collateral channel. This issue is particularly pronounced in the log-log specification (7). A one percent increase in the house price represents a five percent increase in housing net worth for a homeowner at 80% LTV, but only a two percent increase for a homeowner at 50% LTV. Mechanically, there are heterogeneous wealth changes depending on LTV. As a robustness check, we have therefore tried a level specification as well, finding very similar qualitative results. This strengthens the conclusion that the collateral channel is crucial. Third and finally, leverage may be correlated with unobserved individual characteristics that affect borrowing behavior. A candidate would be self-control problems. As [Mian & Sufi \(2011\)](#) note, it is likely that households with greater self-control problems will be observed as credit constrained. However, when augmenting equation (7) with individual fixed effects (which should pick up self-control problems), we find that our heterogeneity results are qualitatively unchanged (albeit with larger standard errors). Controlling for fixed effects ensures that the variation in leverage isn't confounded with individual characteristics, such as self-control.

To conclude, the heterogeneity results are strongly suggestive of the collateral channel, but perhaps not fully conclusive to a sceptic. In the next section, we propose a more sophisticated test — one that exploits discrete changes in the tightness of collateral constraints around interest notches — providing our final piece of evidence in favor of the collateral channel.

5.2 Collateral Channel: A Test Using Interest Notches

The UK setting offers a novel way of investigating the collateral channel arising from the presence of observable credit constraints that depend on collateral. As described in section 2.1, the mortgage interest rate schedule features discrete jumps (notches) at critical LTV thresholds. There are notches

at LTV ratios of 50%, 60%, 70%, 75%, 80%, 85%, and 90%.²² These notches introduce ‘soft’ credit constraints: as the LTV ratio surpasses (and housing collateral therefore falls below) one of the critical thresholds, the cost of borrowing increases sharply.²³ The direct incentive created by these interest notches is for homeowners to choose LTV ratios below one of the thresholds, thus creating bunching in the LTV distribution. Such bunching represents borrowing responses to the interest rate — rather than to the house price — and was studied in detail by *Best et al. (2015)*. Here we consider whether house price movements interact with bunching responses to interest notches in a way that is consistent with the collateral channel. The basic idea is that house price growth, by increasing collateral, moves homeowners below interest notches and induces borrowing responses due to reduced costs of borrowing.

We start by presenting a simple test of heterogeneity, similar in spirit to the preceding analyses, before turning to a more sophisticated and conclusive analysis. Specifically, Figure 14 investigates if the equity extraction elasticity depends on whether the underlying price variation moves homeowners across notches and in which direction. We define the collateral constraint as being relaxed (reinforced) when house price variation moves the homeowner at least one notch down (up) and thus reduces (increases) the mortgage interest rate. Otherwise, the collateral constraint is defined as “unchanged.”²⁴ Panel A of the figure considers a baseline specification without any other controls. This is a specification like (7) in which house price growth is interacted with dummies for the three notch scenarios (relaxed/reinforced/unchanged), but without simultaneously controlling for other dimensions of heterogeneity. This analysis shows that the elasticity is the highest (close to 0.5) when the collateral constraint is relaxed, and that the elasticity is the lowest (close to zero) when the collateral constraint is reinforced.²⁵ The fact that the elasticity is essentially zero when the collateral constraint is reinforced may be due to collateral constraints interacting with liquidity constraints, making it hard for homeowners to inject cash when house price growth increases their cost of borrowing. As a robustness check, Panel B introduces household and month fixed effects in the specification. This graph confirms the qualitative relationship between the bor-

²²Figure A.II illustrates most of these notches.

²³The difference between such soft borrowing constraints and the hard borrowing constraints often assumed in theoretical models can be interpreted in terms of the size of the notch: a hard borrowing constraint is one where the borrowing cost jumps prohibitively at a threshold. In fact, the 90% LTV notch serves as a hard borrowing constraint for most homeowners in our data, because very few lenders have offered mortgage products above this level since the global financial crisis.

²⁴However, this terminology should not be taken too literally: house price appreciation may relax credit constraints even if it does not move homeowners to a lower interest rate notch.

²⁵While the figure pools all years 2005-15, we have checked that the patterns are roughly the same inside and outside the recession years.

rowing elasticity and changes in collateral constraints, although the effect is smaller than in the baseline specification without fixed effects. The asymmetric response to relaxing and tightening borrowing constraints is suggestive of the importance of the collateral channel.

To provide more conclusive evidence of the collateral channel, Figure 15 analyzes the dynamic interaction between house price growth and bunching responses to interest rate notches. This figure focuses on the sample of households who are pulled down to a lower notch by house price growth, i.e. households whose collateral constraint is relaxed. For this analysis, it is useful to formally define three different LTV concepts. First, we define the pre-determined LTV = D_{it}^P / P_{it-1} as the homeowner's LTV at time t given past mortgage choices (the debt level and amortization schedule chosen at time $t - 1$) and the old house price. Second, we define the passive LTV = D_{it}^P / P_{it} as the homeowner's LTV at time t given past mortgage choices and the new house price. This is the LTV that would apply if the homeowner simply rolled over her debt at time t , i.e. if she were "passive." Third, there is the actual chosen LTV = D_{it} / P_{it} that includes any equity extraction or injection at time t . By this terminology, the sample in the figure includes borrowers for whom the passive LTV is at least one notch down from their pre-determined LTV.

The figure shows two panels in which we compare the density distributions of the three LTV measures defined above. The x-axis in each panel represents the distance between a given LTV measure (pre-LTV, passive LTV, or chosen LTV) and the next-notch-up from the passive LTV. Panel A illustrates the implications of house price growth by comparing the distributions of pre-LTV and passive LTV. Two implications are worth highlighting. First, house price growth moves all borrowers from the positive range (in terms of their pre-LTV) to the negative range (in terms of their passive LTV). This follows from the fact that we are restricting the sample to households who are pulled down by at least one notch. Second, house price growth eliminates all bunching at interest notches: there is bunching at every notch in the pre-LTV distribution, but no bunching in the passive LTV distribution.²⁶

How do borrowers respond to the relaxed collateral constraints? Panel B illustrates the implications of equity extraction behavior by comparing the distributions of the passive LTV and the final chosen LTV. Strikingly, equity extraction behavior largely recreates the qualitative pattern that existed before house price growth. We see a dramatic right-shift of the LTV distribution, moving

²⁶The fact that the passive LTV distribution primarily falls in the bins $(-5, 0)$ and $(-10, -5)$, with a discrete drop between the two, is not a bunching response. It follows mechanically from the x-axis normalization and the fact that most homeowners are no longer 5 or 10 percentage points away from a notch. Furthermore, notice that bunching in the pre-LTV distribution is attenuated compared to the actual amount of bunching in the last refinance event due to amortization between the last and current refinance events.

borrowers back to around zero or into the positive range, and recreating bunching at notches. In other words, when house price growth pulls households below one or more notches (Panel A), most of them extract equity back to the next notch above (at zero) or a higher notch (in the positive range). Hence, this figure shows how house price growth interacts with bunching responses to interest notches in a way that is consistent with a collateral mechanism.²⁷

To conclude, the multivariate heterogeneity analysis presented in the previous section combined with the notches analysis presented here provides quite compelling evidence of the importance of the collateral channel. Most previous work on the effect of house prices on household borrowing was only able to estimate the total effect, because the empirical analyses were based on house price variation (regional or aggregate variation) that affect both wealth and collateral.

6 Conclusion

The global financial crisis of 2007-8 has reignited a debate on the role of house prices in driving household debt. A first generation of papers following the crisis studied this question using regional data in the US and found strong borrowing responses. This paper takes a different methodological tack to this question. Rich administrative mortgage data and novel features of the UK mortgage market allow us to construct a large panel of refinancers and study the relationship between borrowing and house prices at the household level. Exploring a far more granular source of house-price variation that relies on idiosyncratic rather than regional variation in house prices, we find household borrowing responses that are considerably smaller than those found for the US using regional data. Specifically, we find that a 10% percent increase in individual house price increases borrowing by 2%. Importantly, when we do not make full use of the panel structure and consider the pre-crisis period, we obtain similar elasticities to those found for the US.

Our rich dataset also allows us to explore why borrowing responds to house prices. The striking finding from this analysis is that there is essentially no heterogeneity in any dimension except one — loan-to-value — but that this dimension is very strong. In particular, the elasticity is strongly increasing in LTV ratios, even after controlling non-parametrically for factors such as borrower age, income and income growth. This heterogeneity analysis together with a test using ‘soft’ credit constraints (interest rate notches based on collateral) strongly suggests that the housing collateral channel is the main driver of the elasticities we find.

²⁷To be clear, what is new in Figure 15 compared to the more standard bunching analysis in Best *et al.* (2015) is the illustration of a dynamic interaction between house price growth and bunching responses.

The magnitude of these responses, and the importance of collateral constraints, has important implications for understanding household behavior in both micro- and macro-economics. A growing literature on macro and housing relies on collateral constraints to obtain realistic macro responses to boom-bust cycles in the housing market. Our findings affirm such theoretical approaches and provide microeconomic estimates that could help discipline future research in this area.

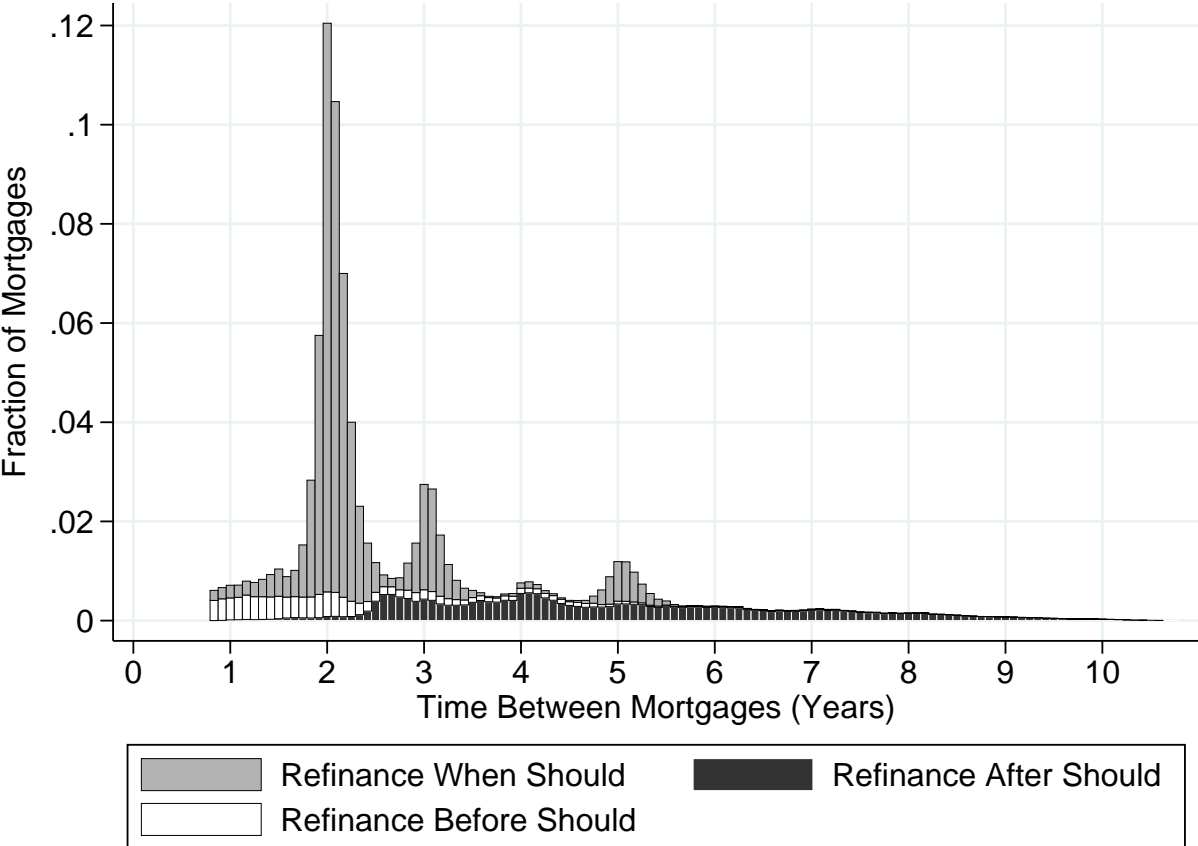
References

- ADELINO, MANUEL, SCHOAR, ANTOINETTE, & SEVERINO, FELIPE. 2015. House Prices, Collateral, and Self-employment. *Journal of Financial Economics*, **117**(2), 288–306. [19](#)
- ADELINO, MANUEL, SCHOAR, ANTOINETTE, & SEVERINO, FELIPE. 2016. Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class. *Review of Financial Studies*, **29**(7), 1635–1670. [1](#)
- AGARWAL, SUMIT, BEN-DAVID, ITZHAK, & YAO, VINCENT W. 2015. Collateral Valuation and Borrower Financial Constraints: Evidence from the Residential Real Estate Market. *Management Science*, **61**(9), 2220–2240. [6](#)
- AGARWAL, SUMIT, AMBROSE, BRENT W., & YAO, VINCENT W. 2016. The Effects and Limits of Regulation: Appraisal Bias in the Mortgage Market. Mimeo. [6](#)
- AOKI, KOSUKE, PROUDMAN, JAMES, & Vlieghe, GERTJAN. 2004. House Prices, Consumption, and Monetary Policy: a Financial Accelerator Approach. *Journal of Financial Intermediation*, **13**(4), 414–435. [19](#)
- ATTANASIO, ORAZIO P., & WEBER, GUGLIELMO. 1994. The UK Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence. *The Economic Journal*, **104**(427), 1269–1302. [19](#)
- ATTANASIO, ORAZIO P., BLOW, LAURA, HAMILTON, ROBERT, & LEICESTER, ANDREW. 2009. Booms and Busts: Consumption, House Prices and Expectations. *Economica*, **76**(301), 20–50. [1, 3, 19, 21](#)
- ATTANASIO, ORAZIO P., LEICESTER, ANDREW, & WAKEFIELD, MATTHEW. 2011. Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK. *Journal of the European Economic Association*, **9**(3), 399–435. [1, 3](#)
- BEN-DAVID, ITZHAK. 2011. Financial Constraints and Inflated Home Prices During the Real Estate Boom. *American Economic Journal: Applied Economics*, **3**(3), 55–87. [6](#)
- BERGER, DAVID, GUERRIERI, VERONICA, LORENZONI, GUIDO, & VAVRA, JOSEPH. 2015. House Prices and Consumer Spending. NBER Working Paper 21667. [2, 19, 21](#)

- BEST, MICHAEL, CLOYNE, JAMES, ILZETZKI, ETHAN, & KLEVEN, HENRIK. 2015. Interest Rates, Debt and Intertemporal Allocation: Evidence From Notched Mortgage Contracts in the UK. Bank of England Staff Working Paper 543. [1](#), [3](#), [6](#), [14](#), [23](#), [25](#), [50](#)
- CAMPBELL, JOHN Y., & COCCO, JOÃO F. 2007. How do House Prices Affect Consumption? Evidence from Micro Data. *Journal of Monetary Economics*, **54**(3), 591–621. [1](#), [3](#), [19](#)
- CASE, KARL E., & SHILLER, ROBERT J. 1989. The Efficiency of the Market for Single-Family Homes. *American Economic Review*, **79**(1), 125–137. [11](#)
- CASE, KARL E., SHILLER, ROBERT J., THOMPSON, ANNE K., LAIBSON, DAVID, & WILLEN, PAUL. 2012. What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets. *Brookings Papers on Economic Activity*, 265–315. [11](#)
- CASE, KARL E., QUIGLEY, JOHN M., & SHILLER, ROBERT J. 2013. Wealth Effects Revisited: 1975–2012. NBER Working Paper 18667. [19](#)
- CHANEY, THOMAS, SRAER, DAVID, & THESMAR, DAVID. 2012. The Collateral Channel: How Real Estate Shocks Affect Corporate Investment. *The American Economic Review*, **102**(6), 2381–2409. [19](#)
- DAVIDOFF, THOMAS. 2013. Supply Elasticity and the Housing Cycle of the 2000s. *Real Estate Economics*, **41**(4), 793–813. [1](#)
- DAVIDOFF, THOMAS. 2016. Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors. *Critical Finance Review*, forthcoming. [1](#)
- DEFUSCO, ANTHONY. 2016. Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls. Mimeo: Kellogg School of Management, Northwestern University. [4](#), [19](#)
- HILBER, CHRISTIAN A.L., & VERMEULEN, WOUTER. 2016. The Impact of Supply Constraints on House Prices in England. *The Economic Journal*, **126**(591), 358–405. [13](#)
- HURST, ERIK, & STAFFORD, FRANK P. 2004. Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption. *Journal of Money, Credit, and Banking*, **36**(6), 985–1014. [2](#), [5](#)
- IACOVIELLO, MATTEO. 2005. House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle. *American Economic Review*, **95**(3), 739–764. [19](#)

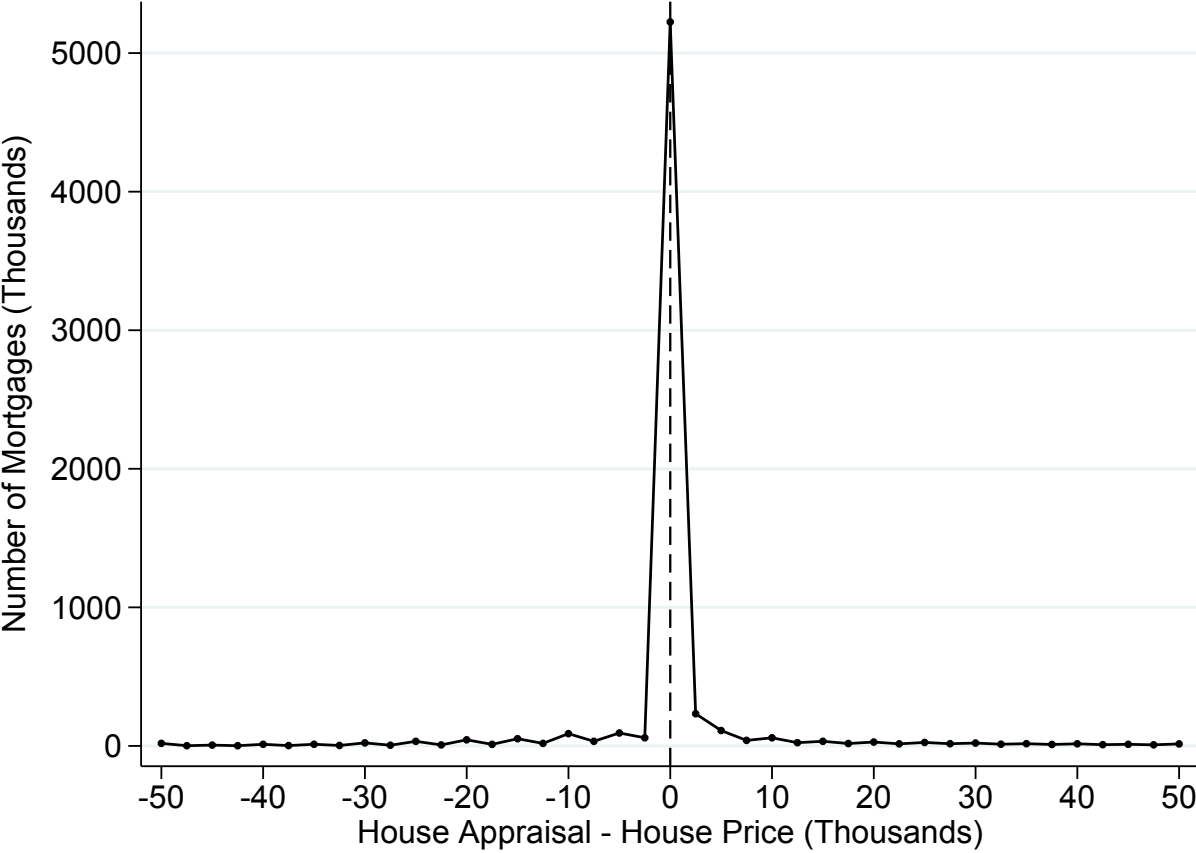
- KAPLAN, GREG, MITMAN, KURT, & VIOLANTE, GIOVANNI L. 2015. Consumption and House Prices in the Great Recession: Model Meets Evidence. Mimeo. [1](#), [4](#), [11](#)
- KLEVEN, HENRIK. 2016. Bunching. *Annual Review of Economics*, **8**, 435–464. [9](#)
- LUSTIG, HANNO N., & NIEUWERBURGH, STIJN G. VAN. 2005. Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective. *Journal of Finance*, **60**(3), 1167–1219. [19](#)
- MIAN, ATIF, & SUFI, AMIR. 2009. The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *Quarterly Journal of Economics*, **124**(4), 1449–1496. [3](#)
- MIAN, ATIF, & SUFI, AMIR. 2011. House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis. *American Economic Review*, **101**(5), 2132–2156. [1](#), [2](#), [3](#), [4](#), [13](#), [19](#), [22](#)
- MIAN, ATIF, & SUFI, AMIR. 2014. *House of Debt: How They (And You) Caused the Great Recession, and How We Can Prevent It From Happening Again*. University of Chicago Press. [1](#)
- MIAN, ATIF, RAO, KAMALESH, & SUFI, AMIR. 2013. Household Balance Sheets, Consumption, and the Economic Slump. *The Quarterly Journal of Economics*, **128**(4), 1687–1726. [1](#)
- SAIZ, ALBERT. 2010. The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics*, **125**(3), 1253–1296. [1](#), [13](#)
- SHILLER, ROBERT J. 2007. Understanding Recent Trends in House Prices and Home Ownership. *Proceedings - Economic Policy Symposium - Jackson Hole, Federal Reserve Bank of Kansas City*, 89–123. [11](#)
- SINAI, TODD, & SOULELES, NICHOLAS S. 2005. Owner-Occupied Housing as a Hedge Against Rent Risk. *The Quarterly Journal of Economics*, **120**(2), 763–789. [2](#), [19](#)

FIGURE 1: HOMEOWNERS REFINANCE AROUND THE ONSET OF THE RESET RATE



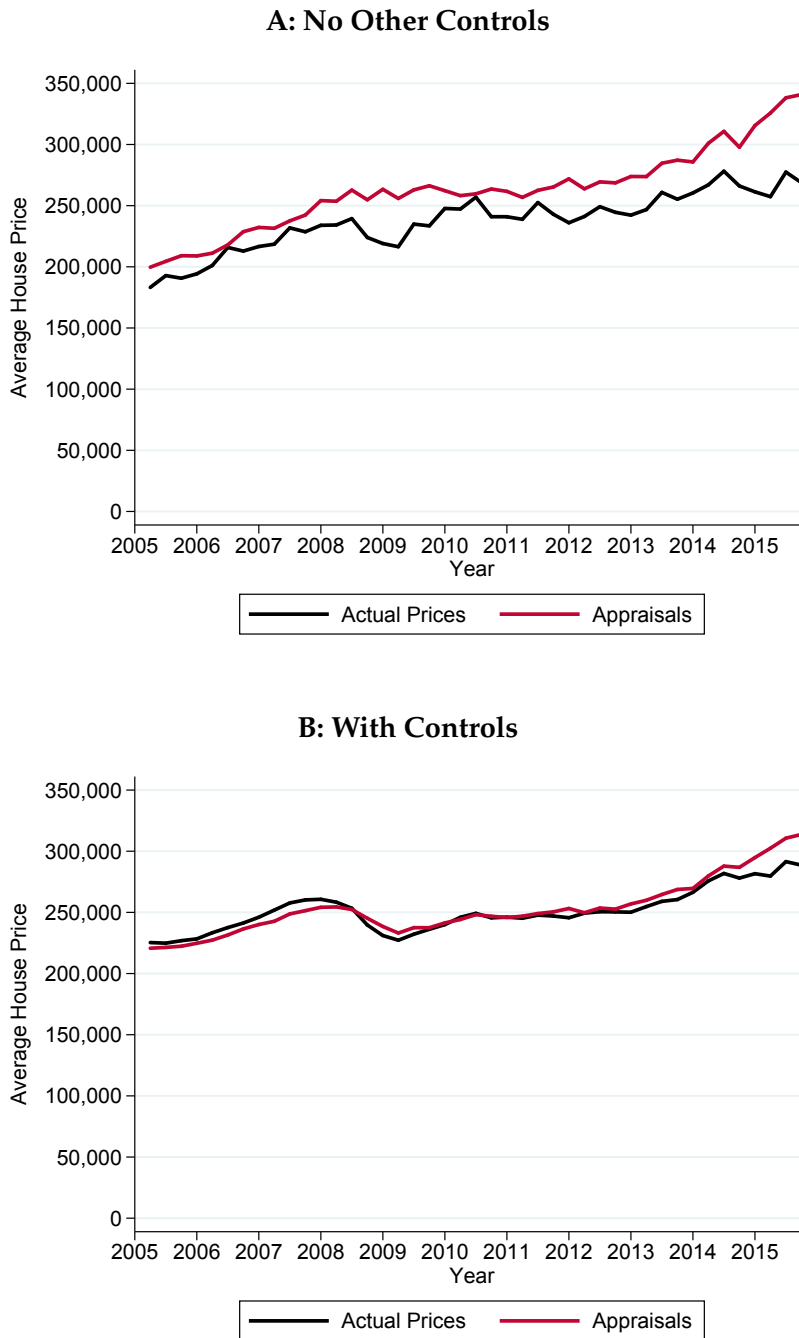
Notes: The figure shows the distribution of the time between mortgage financing events. Households who refinance between 2 months before and 6 months after the onset of their reset rate are shown in light gray, households who refinance more than 6 months after the onset of their reset rate are shown in black, and households who refinance more than 2 months before the onset of their reset rate are shown in white. The data in this figure exclude households for whom we do not observe the date of reset rate onset.

FIGURE 2: HOUSE PRICES VS APPRAISALS (NEW PURCHASES)



Notes: The figure shows the distribution of within-house differences between the actual house price and the appraisal price for transacted properties. This includes both first time buyers and home movers, but not refinancers.

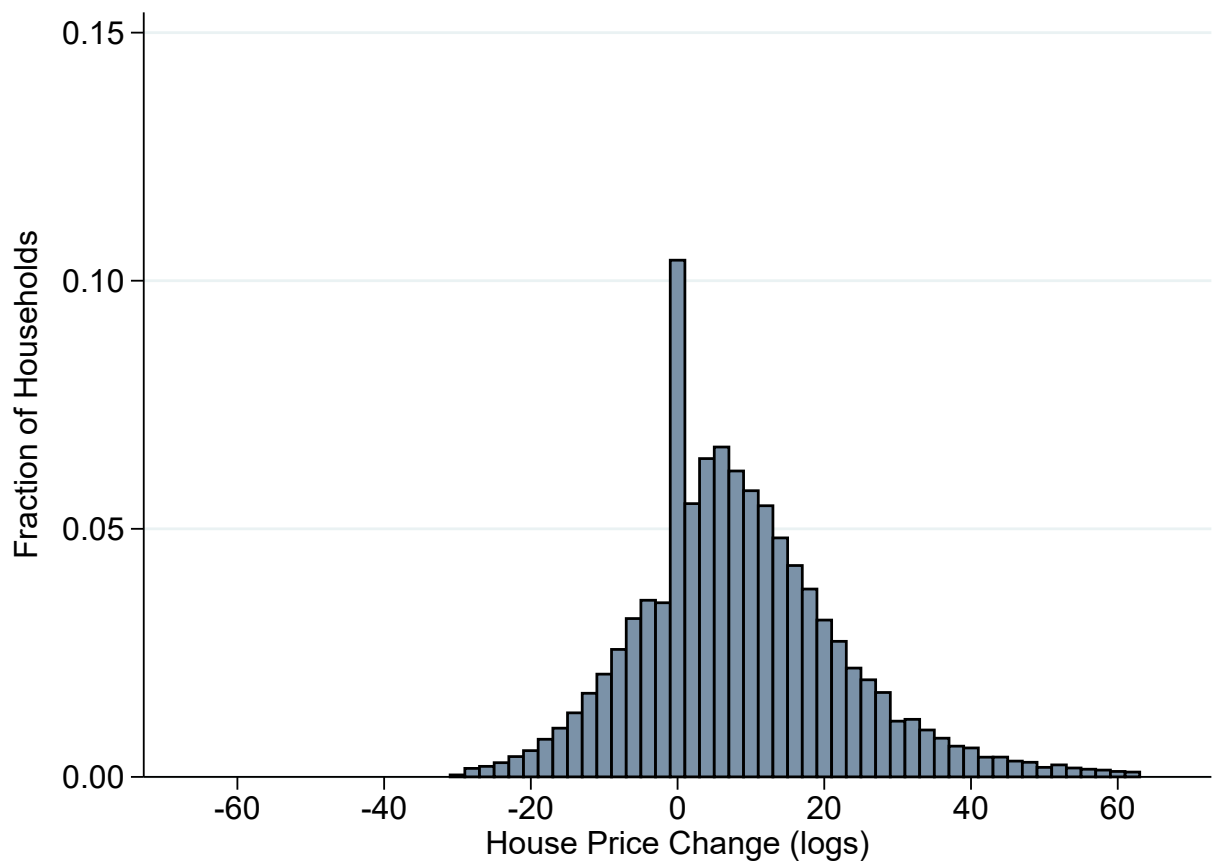
FIGURE 3: HOUSE PRICES VS APPRAISALS (REFINANCED HOMES)



Notes: The figure compares actual house prices (for transacted properties) with appraisal prices (for refinanced properties) over time. Panel A plots the raw time series of actual and appraised prices, obtained by regressing each of the price series on a full set of quarter dummies and plotting the estimated coefficients. Panel B augments the price regressions on quarter dummies with controls for twenty quantiles of the borrower age distribution as well as twenty quantiles of the postcode-level price distribution (see equation (1)). The panel plots the coefficients on the quarter dummies, plus a constant equal to the effect of the average age and the average postcode. This panel shows that, once we correct for compositional differences in age and postcode, there is no significant appraisal bias.

FIGURE 4: DISTRIBUTION OF RAW HOUSE PRICE GROWTH

Households With ≥ 2 Mortgage Observations

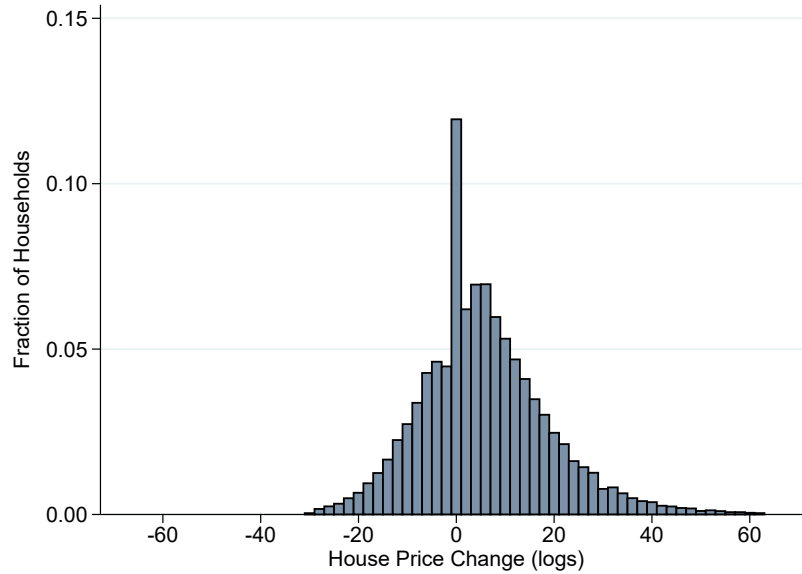


Notes: The figure shows the distribution of raw house price growth among households for whom we observe at least two mortgage financing events. House price growth is measured as the log change in house prices between the current and the last mortgage event, multiplied by 100 (i.e., approximately percentage house price growth).

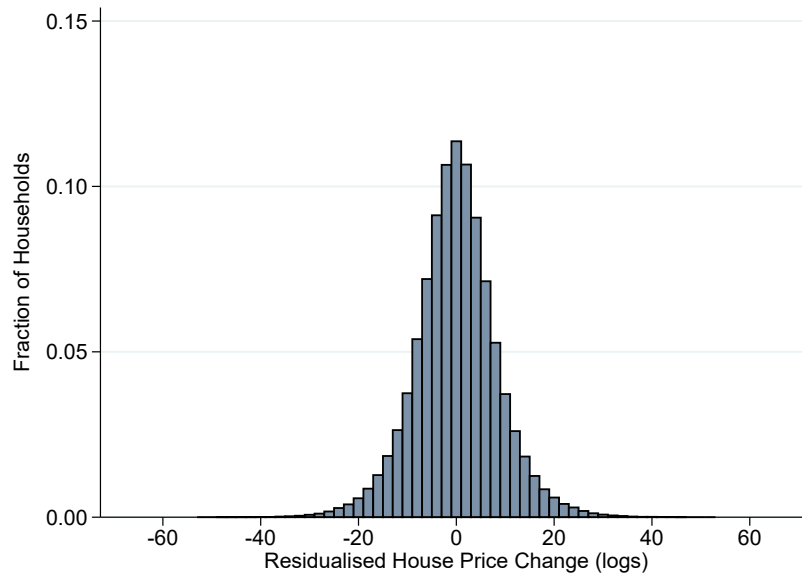
FIGURE 5: DISTRIBUTION OF RAW VS RESIDUALIZED HOUSE PRICE GROWTH

Households With ≥ 3 Mortgage Observations

A: Raw Price Growth



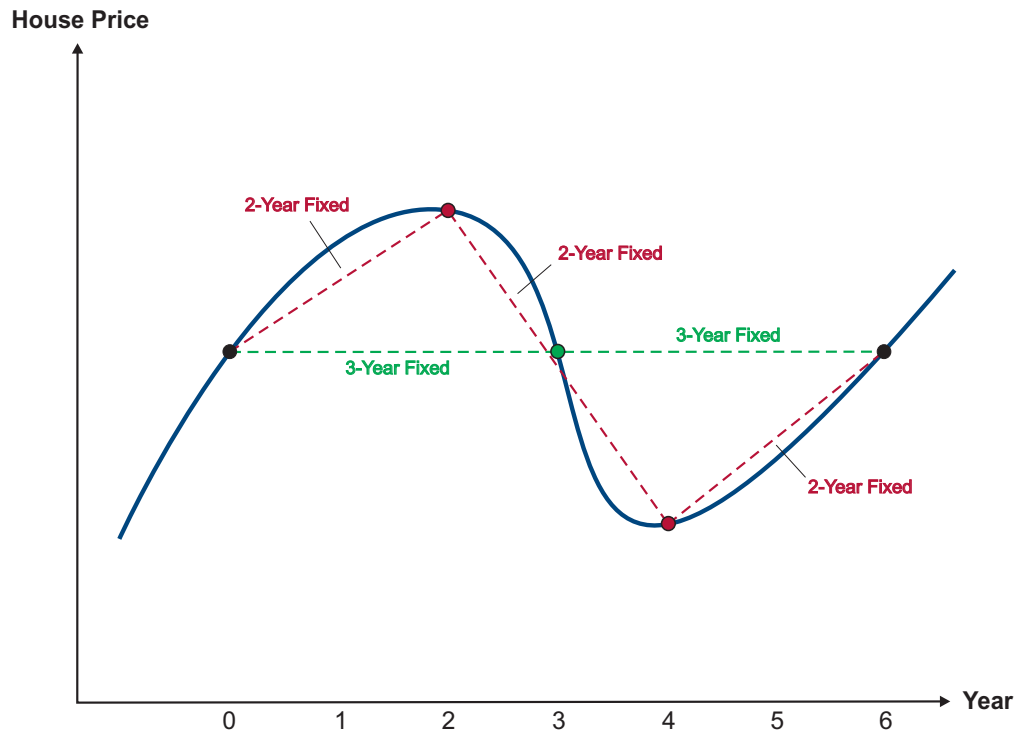
B: Residualized Price Growth After Absorbing Fixed Effects



Notes: The figure shows distributions of house price growth among households for whom we observe at least three mortgage financing events. Panel A shows the distribution of raw house price growth, while Panel B shows the distribution of residualized house price growth after absorbing household fixed effects, month fixed effects, and county-by-year fixed effects. In both panels, house price growth is measured as the log change in house prices between the current and the last mortgage event, multiplied by 100 (i.e., approximately percentage house price growth).

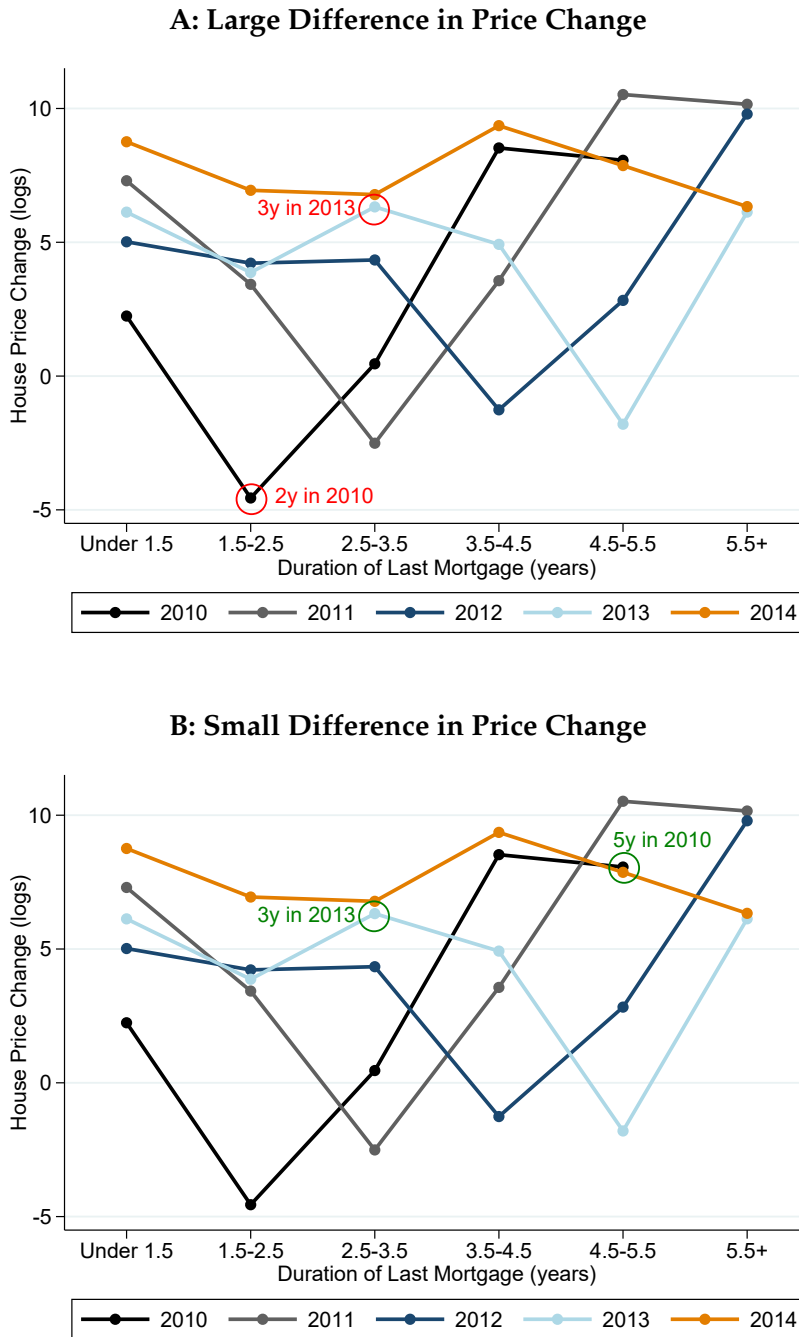
FIGURE 6: THE TIMING OF REFINANCE EVENTS AND HOUSE PRICE CHANGES

Conceptual Example



Notes: The figure illustrates, in a conceptual example, how differences in contract duration choices create variation in house price changes across households. The graph compares two homeowners who start out at the same time (time 0), live in houses with the same price cycle (the solid blue line), but have different preferences over low-interest rate durations. One homeowner prefers 2-year fixed interest rate loans, while the other prefers 3-year fixed interest loans. The homeowner in two-year contracts refinances three times over a 6-year period, facing either positive or negative price shocks at each event, whereas the homeowner in 3-year contracts refinances only two times facing a zero price shock each time. Our empirical strategy exploits such within-person variation in price growth driven by duration choices.

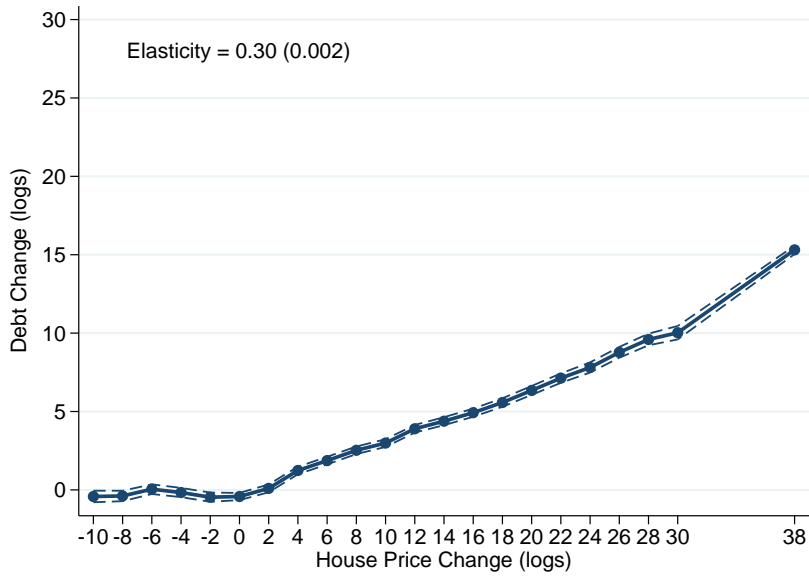
FIGURE 7: HOUSE PRICE CHANGES VS LAST DURATION X TIME OF REFINANCE



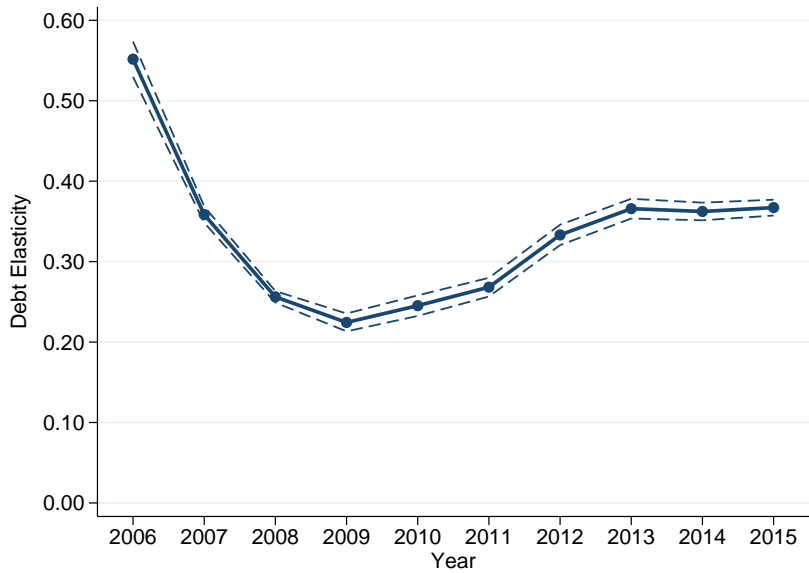
Notes: This figure is the empirical counterpart to the preceding conceptual figure. It plots average house price growth between refinance events for homeowners who refinance at different points in time (in January of different years) by bins of the duration of their last mortgage (number of years between the current and the last refinance events). The two panels show the same graphs, but highlight two different homeowners who experience very different within-person price patterns due to past duration choices. The homeowner in Panel A experiences a large negative price change in January 2010, followed by a large positive change in January 2013. The homeowner in Panel B also refinances in January 2010 and January 2013, but experiences similar price changes in the two events. Our empirical approach uses such within-person variation for identification.

FIGURE 8: MORTGAGE DEBT AND HOUSE PRICES

A: Debt Growth vs House Price Growth



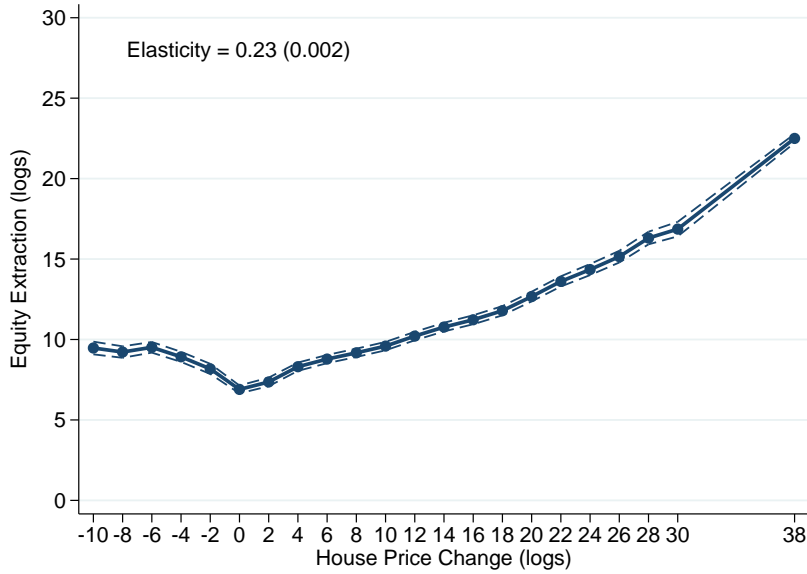
B: Debt Elasticity by Year



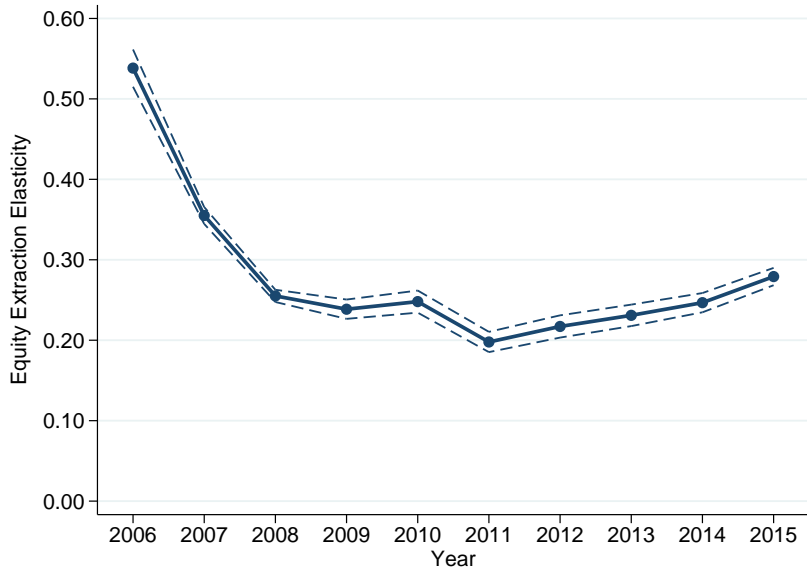
Notes: Panel A plots the average mortgage debt growth in different bins of house price growth, pooling all years 2005-15. Debt growth and house price growth are measured as log changes between refinance events multiplied by 100 (i.e., approximately percentage changes). The dashed lines represent 95% confidence intervals based on standard errors clustered at the household level. Panel A also shows the average elasticity of mortgage debt with respect to house prices across all years. Panel B reports the elasticity for each year separately, showing that the elasticity is pro-cyclical.

FIGURE 9: EQUITY EXTRACTION AND HOUSE PRICES

A: Equity Extraction vs House Price Growth

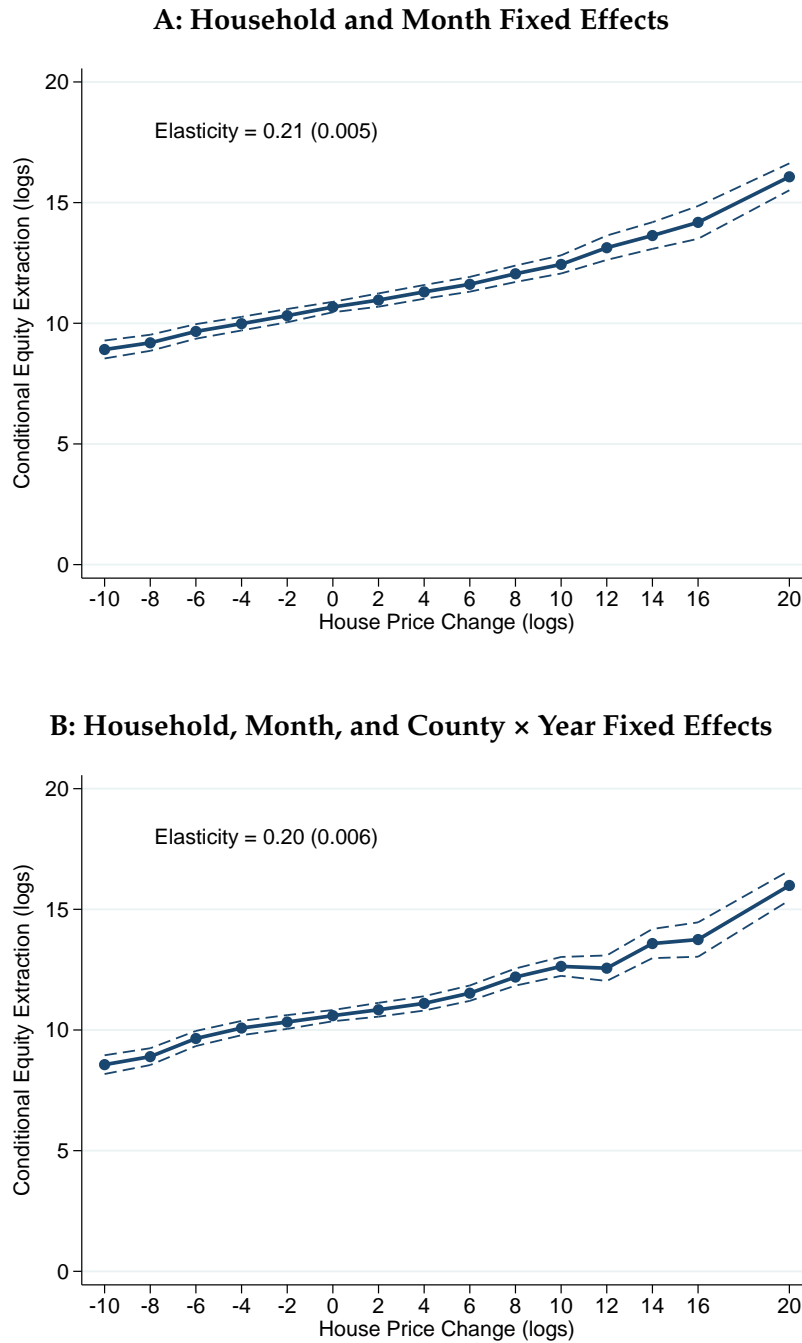


B: Equity Extraction Elasticity by Year



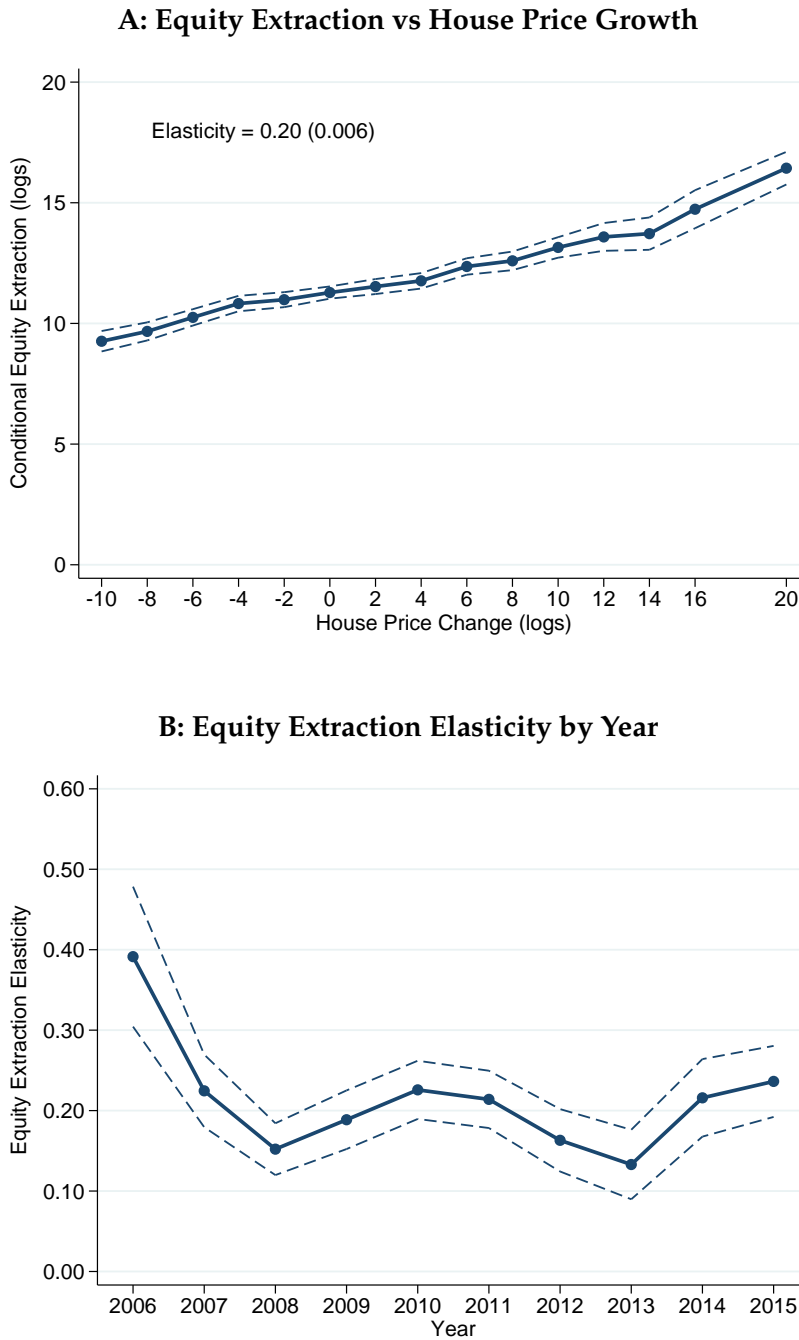
Notes: This figure corresponds to the previous figure, but considers equity extraction when refinancing (as opposed to total debt growth between refinance events) as the outcome variable. Equity extraction is measured as the log difference between mortgage debt after refinancing and the outstanding mortgage debt just before refinancing (i.e., the debt the household would hold if she simply rolled over the existing mortgage debt at the time of refinancing), multiplied by 100. Panel A plots average equity extraction in different bins of house price growth, pooling all years 2005-15. The dashed lines represent 95% confidence intervals, with standard errors clustered by household. Panel A also reports the average equity extraction elasticity across all years, while Panel B shows the equity extraction elasticity for each year separately.

FIGURE 10: EQUITY EXTRACTION AND HOUSE PRICES WITH FIXED EFFECTS



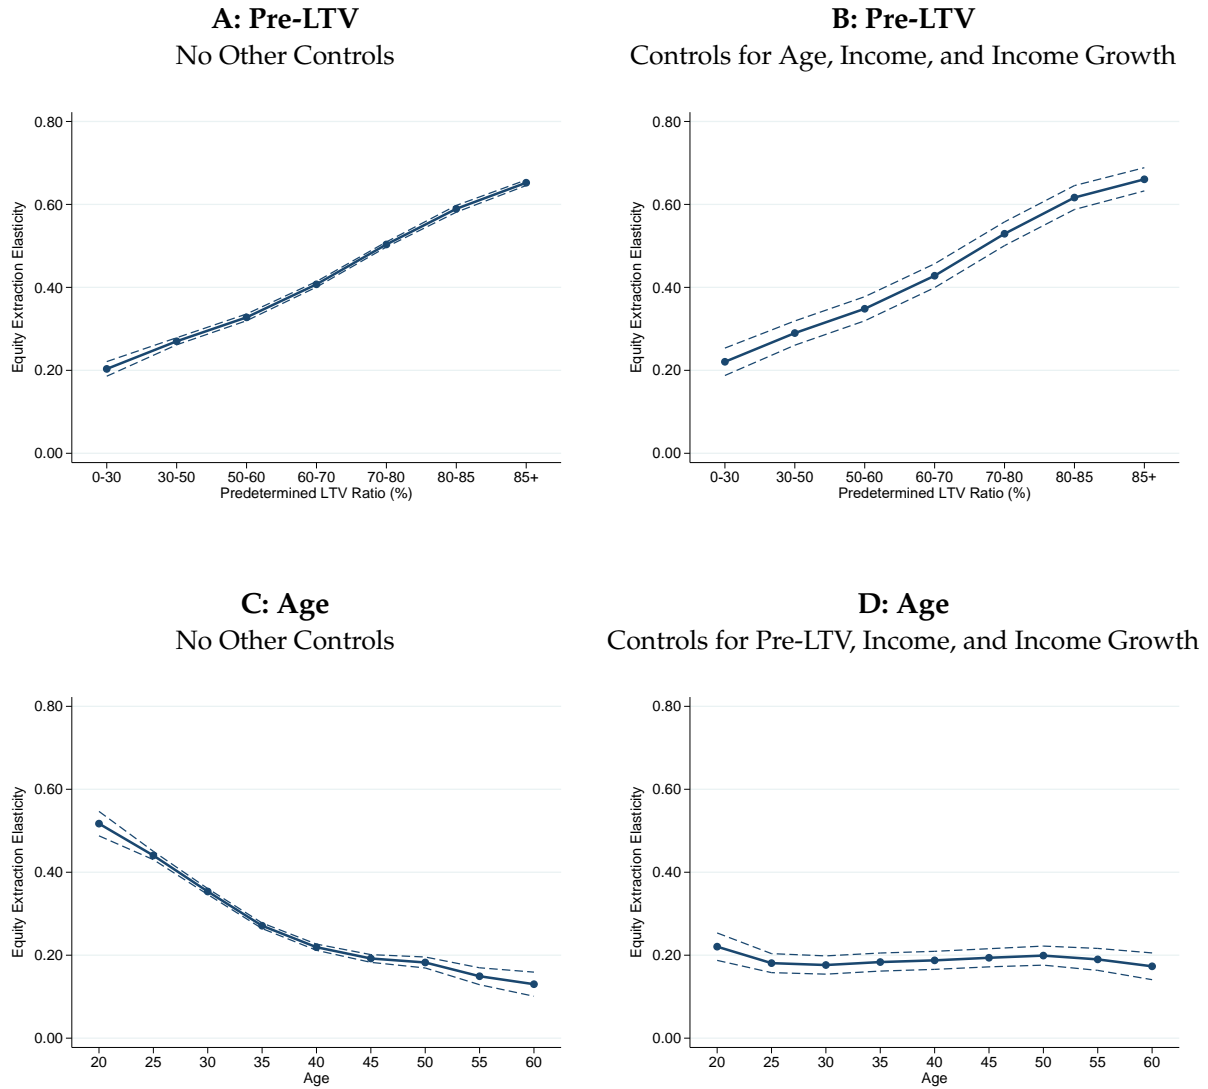
Notes: These panels plot conditional equity extraction in different bins of house price growth based on the fixed effects specification (3), pooling all years 2005-15. The plotted points are the estimated coefficients on house price growth dummies, adding a constant equal to the mean predicted value of equity extraction from all the other covariates. In Panel A, the other covariates are fixed effects for household and month. In Panel B, the other covariates are fixed effects for household, month, and county \times year. The dashed lines represent 95% confidence intervals based on standard errors clustered by household. Each panel reports the average equity extraction elasticity based on a log-linear specification. The figure shows an almost perfectly log-linear relationship between equity extraction and house prices, and it shows that the relationship is unaffected by county \times year (conditional on the other fixed effects).

FIGURE 11: EQUITY EXTRACTION AND HOUSE PRICES WITH FIXED EFFECTS AND HOUSEHOLD CONTROLS



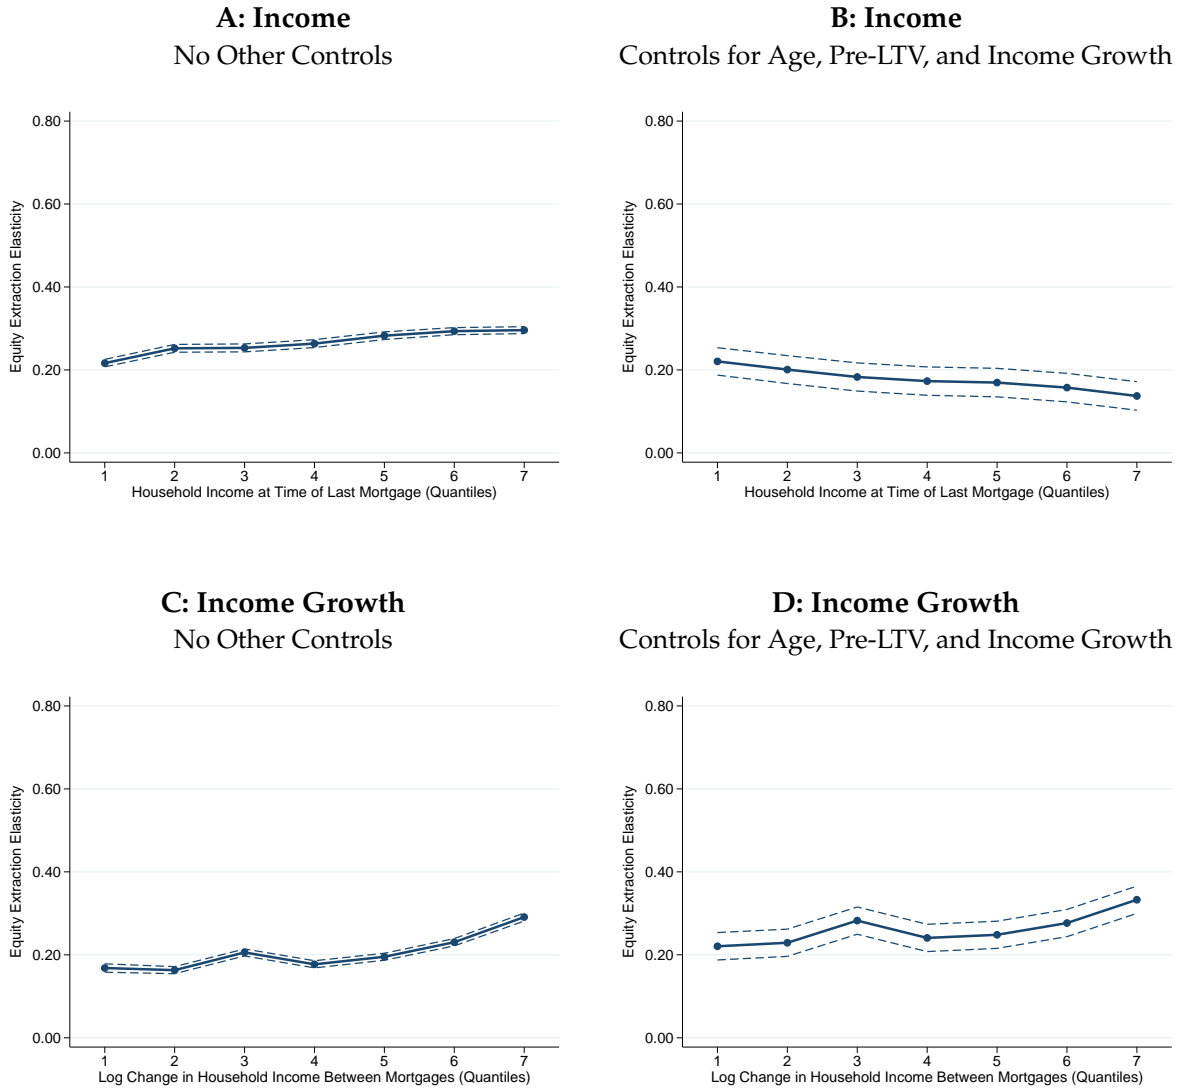
Notes: Panel A plots conditional equity extraction in different bins of house price growth based on the specification with fixed effects and household-level controls in equation (4). The panel is constructed exactly like the previous figure that is based on specifications without household-level controls. The household controls included here are income level, income growth, mortgage interest rate, borrower age, a dummy for couples, and dummies for a range of self-reported reasons for the current and the last refinances. The figure shows that the inclusion of such rich controls makes no difference to the results. While Panel A is based on all years 2005-15, Panel B shows the equity extraction elasticity for each year separately. The rich specification considered here has reduced, but not eliminated, the cyclical in the equity extraction elasticity.

FIGURE 12: HETEROGENEITY IN BORROWING ELASTICITY BY LTV AND AGE



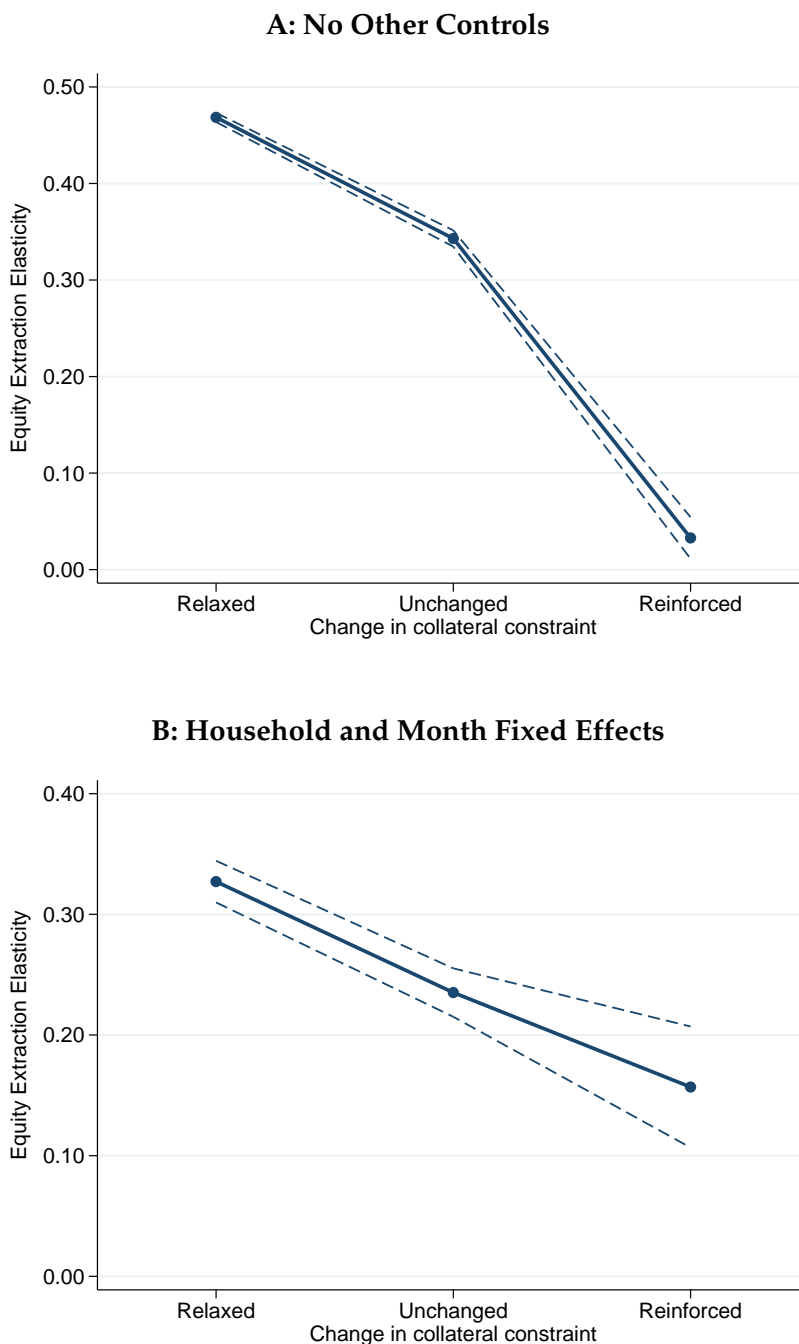
Notes: The figure shows heterogeneity in the equity extraction elasticity by LTV (top panels) and by borrower age (bottom panels). The heterogeneity analysis is based on a pre-determined LTV ratio, namely the LTV ratio at time t absent any equity extraction/injection at time t and absent any house price growth between t and $t - 1$. The left panels are based on univariate specifications that investigate each dimension of heterogeneity on its own, while the right panels are based on multivariate specifications allowing for heterogeneity in four dimensions simultaneously: LTV, borrower age, income level, and income growth. The multivariate specification is shown in equation (7). The dashed lines give 95% confidence intervals based on standard errors clustered by household. The top panels show a strong increasing relationship between LTV and the borrowing elasticity, consistent with collateral effects. The bottom panels show a negative or flat relationship between age and the borrowing elasticity, inconsistent with wealth effects.

FIGURE 13: HETEROGENEITY IN BORROWING ELASTICITY BY INCOME AND INCOME GROWTH



Notes: The figure shows heterogeneity in the equity extraction elasticity by income level (top panels) and by income growth (bottom panels). The income level is measured at the time of the last refinance event, while income growth is measured as the log-change since the last refinance event. The left panels are based on univariate specifications that investigate each dimension of heterogeneity on its own, while the right panels are based on multivariate specifications allowing for heterogeneity in four dimensions simultaneously: LTV, borrower age, income level, and income growth. The multivariate specification is shown in equation (7). The dashed lines give 95% confidence intervals based on standard errors clustered by household. The figure shows that there is relatively little heterogeneity in the borrowing elasticity by either income level or income growth.

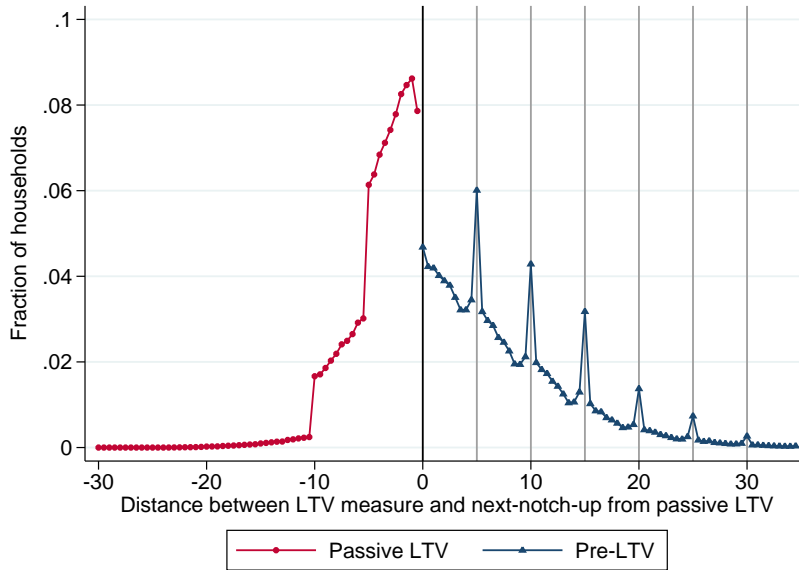
FIGURE 14: HETEROGENEITY IN BORROWING ELASTICITY BY NOTCHES MOVED



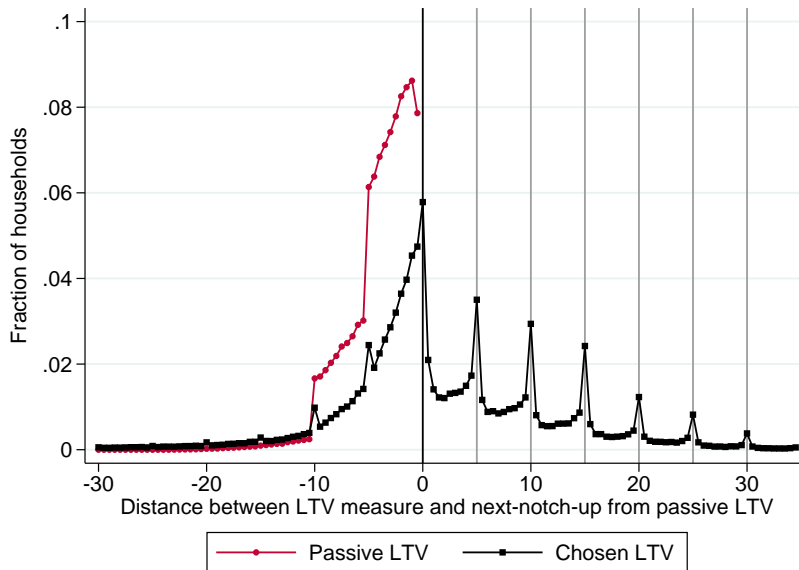
Notes: The figure shows heterogeneity in the equity extraction elasticity by notches moved due to house price changes. There are interest rate notches at LTV thresholds of 50%, 60%, 70%, 75%, 80%, 85%, and 90%. We define the collateral constraint as being relaxed (reinforced) when house price variation moves the homeowner at least one notch down (up) and thus reduces (increases) the interest rate on borrowing. Otherwise, the collateral constraint is defined as “unchanged.” Panel A shows elasticity estimates when including no other controls, while Panel B allows for household and month fixed effects. The dashed lines give 95% confidence intervals based on standard errors clustered by household.

FIGURE 15: HOUSE PRICE GROWTH AND BUNCHING AT COLLATERAL NOTCHES

A: Effect of House Price Growth on LTV Distribution



B: Effect of Equity Extraction on LTV Distribution



Notes: The figure is based on a sample of households who are pulled down to a lower notch by house price growth. The two panels show density distributions of three different LTV measures. The pre-LTV = D_{it}^P / P_{it-1} is the homeowner's LTV at time t given past mortgage choices (i.e., the debt level and amortization schedule chosen at time $t - 1$, not including equity extraction at time t) and the old house price. The passive LTV = D_{it}^P / P_{it} is the homeowner's LTV given past mortgage choices and the new house price. The chosen LTV = D_{it} / P_{it} includes any equity extraction made at time t . The x-axis in each panel represents the distance between a given LTV measure and the next-notch-up from the passive LTV. Panel A illustrates the effects of house price growth by comparing the distributions of pre-LTV and passive LTV. This panel shows that house price growth moves homeowners from the positive to the negative range and eliminates bunching at interest rate notches. Panel B illustrates the effects of borrowing responses by comparing the distributions of the passive LTV and the chosen LTV. This panel shows that equity extraction largely recreates the qualitative pattern that existed before house price growth.

TABLE 1: DESCRIPTIVE STATISTICS

	(1)	(2)	(3)
Panel A: Buyers vs Refinancers			
	Buyers	All Refinancers	Refinancers in our Estimation Sample
Borrower Age	36.47 (10.13)	42.08 (9.77)	40.85 (8.90)
Couple	0.52 (0.50)	0.54 (0.50)	0.54 (0.50)
Income	55,282 (556,583)	55,950 (145,816)	57,603 (81,441)
Income Change (logs)		0.08 (0.36)	0.08 (0.35)
Interest Rate	4.39 (1.40)	4.51 (1.40)	3.98 (1.50)
House Price	229,375 (326,210)	248,329 (361,736)	256,517 (187,020)
LTV	70.72 (21.67)	56.53 (21.80)	61.50 (18.96)
Observations	7,119,807	5,935,441	1,384,346
Panel B: Refinancers in our Estimation Sample			
	Refinance On-Time	Refinance Off-Time	Missing Duration
Borrower Age	39.77 (8.69)	41.58 (8.79)	41.37 (9.04)
Couple	0.55 (0.50)	0.53 (0.50)	0.54 (0.50)
Income	54,516 (48,424)	53,443 (52,356)	62,006 (108,734)
Income Change (logs)	0.08 (0.31)	0.11 (0.38)	0.07 (0.37)
Interest Rate	4.22 (1.51)	3.60 (1.33)	3.97 (1.53)
House Price	245,031 (163,128)	233,110 (158,359)	276,638 (213,290)
LTV	61.56 (18.30)	63.04 (19.27)	60.72 (19.27)
Observations	483,852	288,578	611,916

Notes: The table reports means and standard deviations (in parentheses) for different samples. Panel A compares descriptive statistics for home buyers (column 1), all refinancers (column 2), and refinancers in our estimation sample (column 3). Our estimation sample includes homeowners who we observe refinancing at least twice, and for whom we have sufficient information to precisely measure equity extraction. Panel B compares descriptive statistics for three subsamples of the refinancers in our estimation sample: households who refinance on-time (between two months before and six months after the onset of their reset rate), households who refinance off-time, and households where we do not observe the exact onset of the reset rate.

TABLE 2: EQUITY EXTRACTION ELASTICITIES BY REFINANCE TIMING

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Equity Extraction Elasticity	0.234 (0.002)	0.208 (0.005)	0.204 (0.006)	0.197 (0.006)
Observations	1,384,346	1,384,346	1,311,734	1,173,626
Panel B: On-Time Sample				
Equity Extraction Elasticity	0.245 (0.002)	0.183 (0.006)	0.175 (0.007)	0.166 (0.007)
Observations	483,852	483,852	460,077	459,571
Panel C: Off-Time Sample				
Equity Extraction Elasticity	0.317 (0.004)	0.269 (0.011)	0.263 (0.012)	0.252 (0.013)
Observations	288,578	288,578	274,600	273,727
Panel D: Sample With Missing Durations				
Equity Extraction Elasticity	0.188 (0.003)	0.201 (0.007)	0.202 (0.007)	0.197 (0.009)
Observations	611,916	611,916	577,057	440,328
Control Variables:				
Month FE		×	×	×
Household FE		×	×	×
County x Year FE			×	×
Household Controls				×

Notes: The table reports estimates of the equity extraction elasticity across different specifications and samples. Panel A considers the full sample (summarizing the results of the preceding figures), panel B considers the sample of on-time refinancers (defined as those who refinance between 2 months before and 6 months after reset rate onset), panel C considers the sample of off-time refinancers (defined as those who refinance more than 2 months before or more than 6 months after reset rate onset), and panel D considers the sample of refinancers with missing duration information. Standard errors are clustered by household and shown in parentheses. The household controls included in column (4) are income level, income growth, the last mortgage interest rate, the age of the borrower, a dummy for couples, and dummies for a range of self-reported reasons for the current and the last refinance (pure refinance / home improvement / debt consolidation / other).

TABLE 3: EQUITY EXTRACTION ELASTICITIES USING INSTRUMENTAL VARIABLES

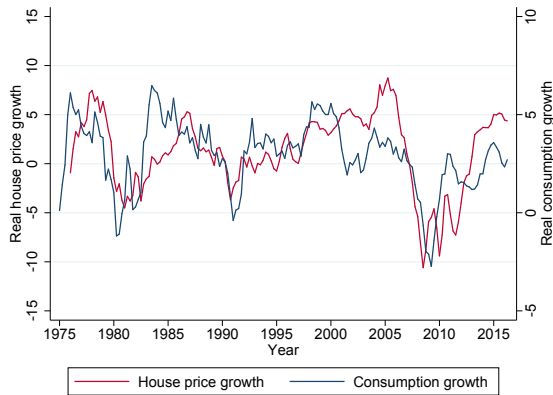
	(1)	(2)	(3)	(4)	(5)
IV Equity Extraction Elasticity	0.150 (0.004)	0.163 (0.004)	0.284 (0.026)	0.295 (0.056)	0.283 (0.056)
Observations	772,430	772,430	772,430	737,168	733,614
<u>Control Variables:</u>					
Contract Duration FE		×	×	×	×
Month FE			×	×	×
Household FE			×	×	×
County x Year FE				×	×
Household Controls					×

Notes: The table reports estimates of the equity extraction elasticity using instrumental variables. The instruments are interactions of dummies for the last mortgage contract duration (time until reset), year and region. The table shows IV elasticities from five different specifications, with the richest specification in column (5) corresponding to equations (5)-(6). The household controls included in column (5) are income level, income growth, the last mortgage interest rate, the age of the borrower, a dummy for couples, and dummies for a range of self-reported reasons for the current and the last refinance (pure refinance / home improvement / debt consolidation / other). Standard errors are clustered by household and given in parentheses. The IV elasticities with fixed effects (equal to 0.28-0.29) are slightly higher than the OLS elasticities, but overall the table confirms the qualitative results.

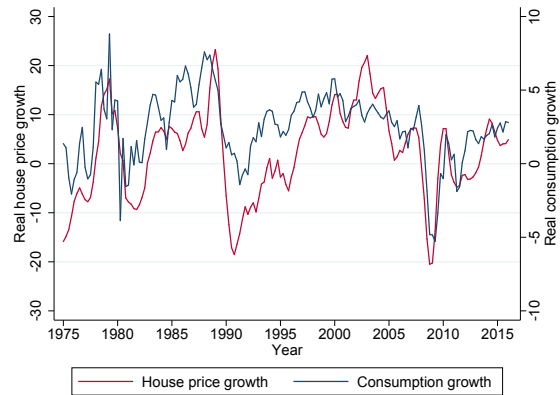
A Appendix

FIGURE A.I: AGGREGATE HOUSE PRICES, CONSUMPTION, AND MORTGAGE DEBT

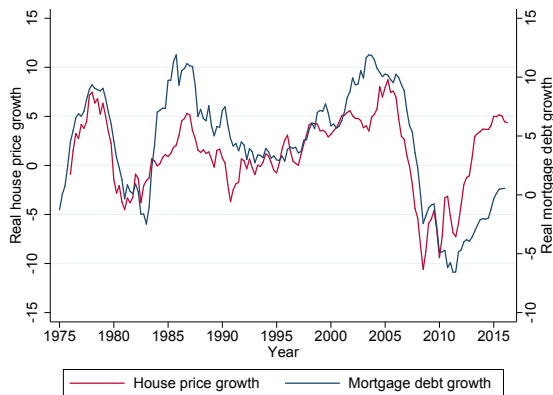
A: US House Price vs Consumption Growth



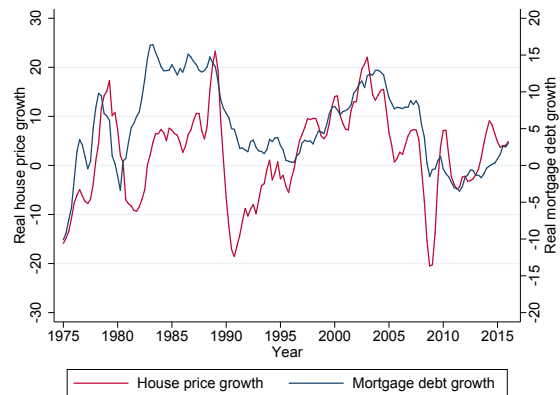
B: UK House Price vs Consumption Growth



C: US House Price vs Mortgage Debt Growth

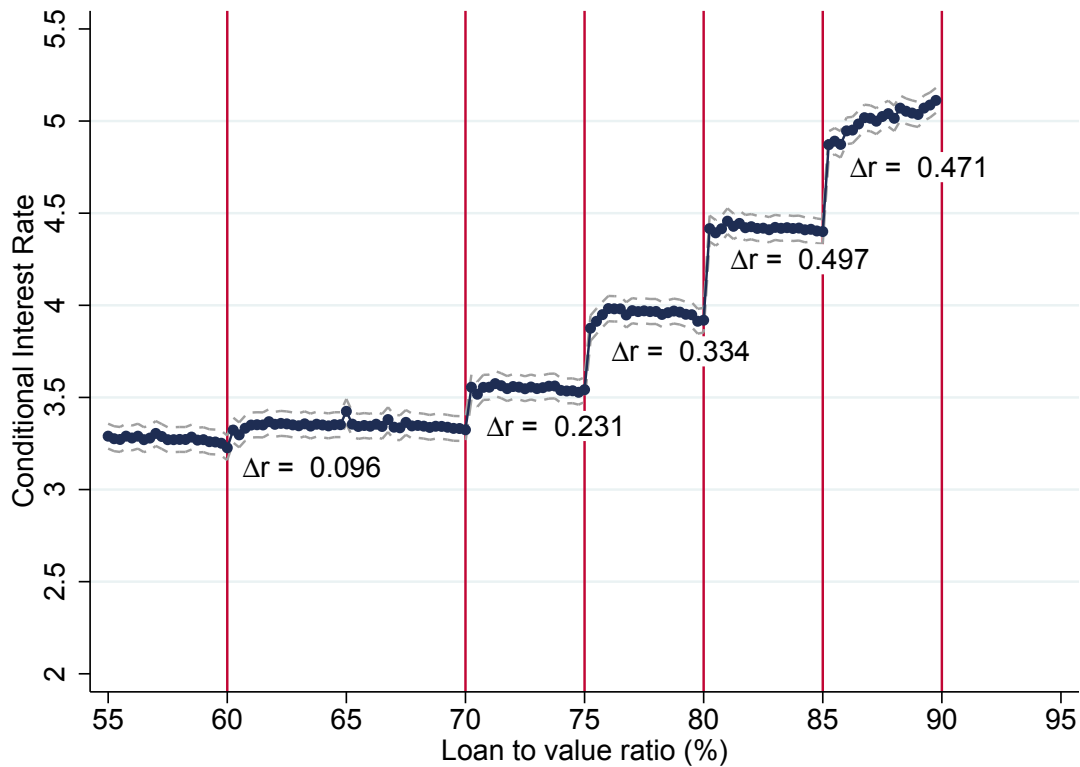


D: UK House Price vs Mortgage Debt Growth



Notes: US house price data are from the Federal Reserve Economic Data, US consumption data are from the BEA National Income and Product Accounts, and US mortgage debt data are from the US Flow of Funds. UK house price data are from the Nationwide Index, UK consumption data are from the ONS National Accounts, and UK mortgage debt data are from the Bank of England. All growth rates are log differences multiplied by 100.

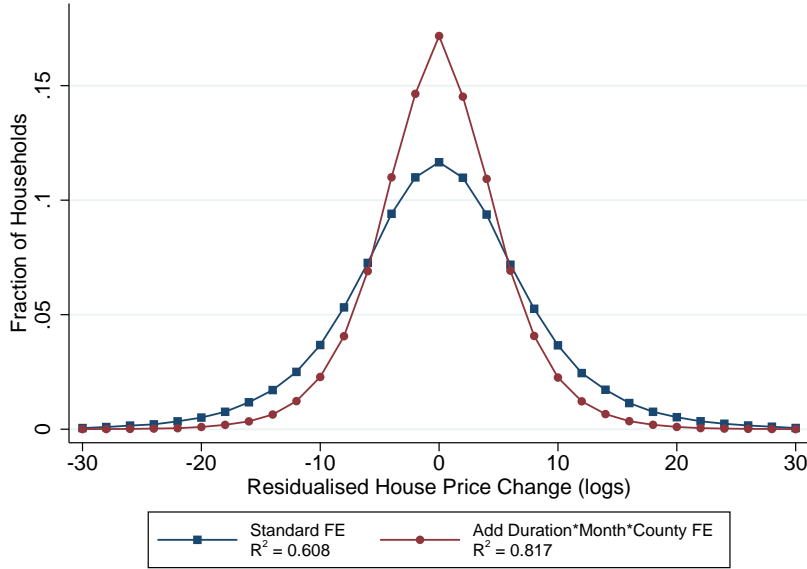
FIGURE A.II: AVERAGE INTEREST RATE SCHEDULE IN THE UK (NOTCHES)



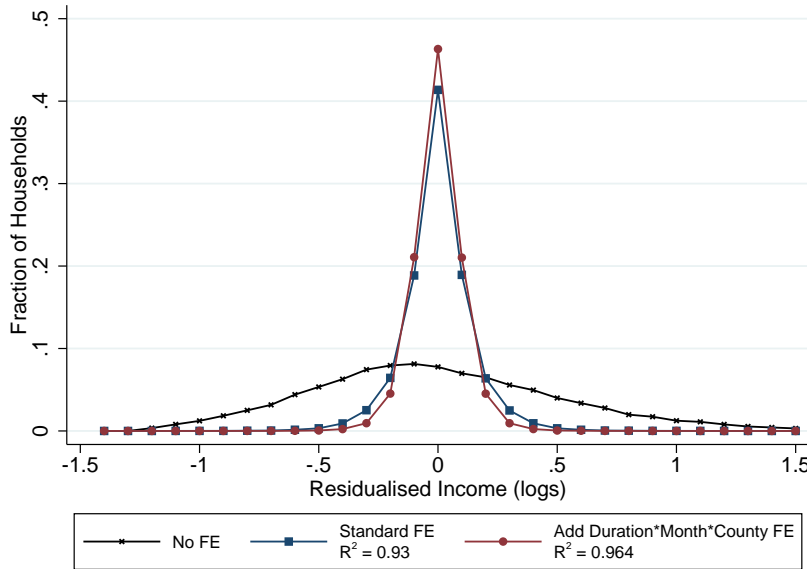
Notes: The figure shows the average mortgage interest rate in the UK (in %) as a step function of the LTV ratio, with sharp jumps (notches) at LTVs of 60%, 70%, 75%, 80%, and 85%. The figure plots coefficients (and confidence intervals) from a regression of the mortgage interest rate on dummies for each 0.25%-bin of the LTV distribution. To each coefficient, we add a constant equal to the mean predicted value of the interest rate from all the other covariates. The other covariates include non-parametric controls for lender, contract duration (time until reset), month of refinance, mortgage type (fixed interest rate / variable interest rate / capped interest rate / other), repayment type (interest only / capital and interest / other), term length, reason for refinance, borrower age, couple indicator, and income. The figure is taken from [Best et al. \(2015\)](#) and further details are provided there.

FIGURE A.III: THE EXPLANATORY POWER OF MORTGAGE DURATION

A: Duration Explains Future Price Change



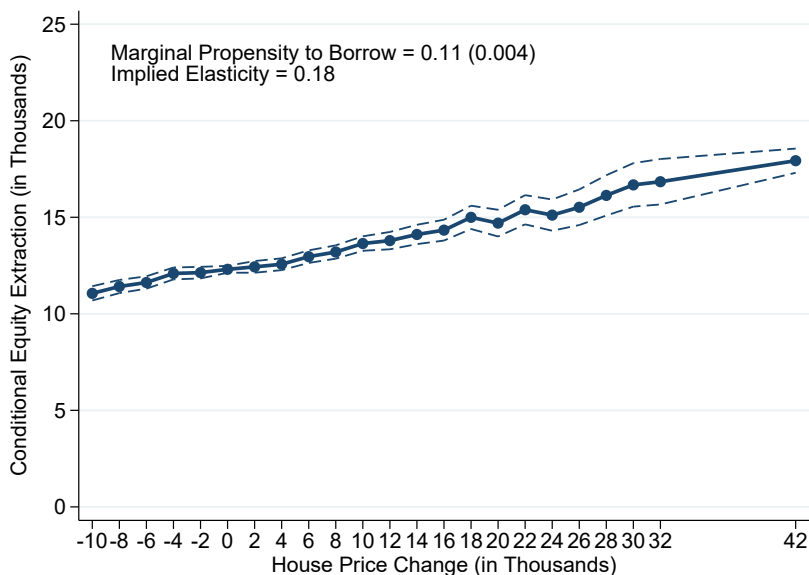
B: Duration Does Not Explain Future Income



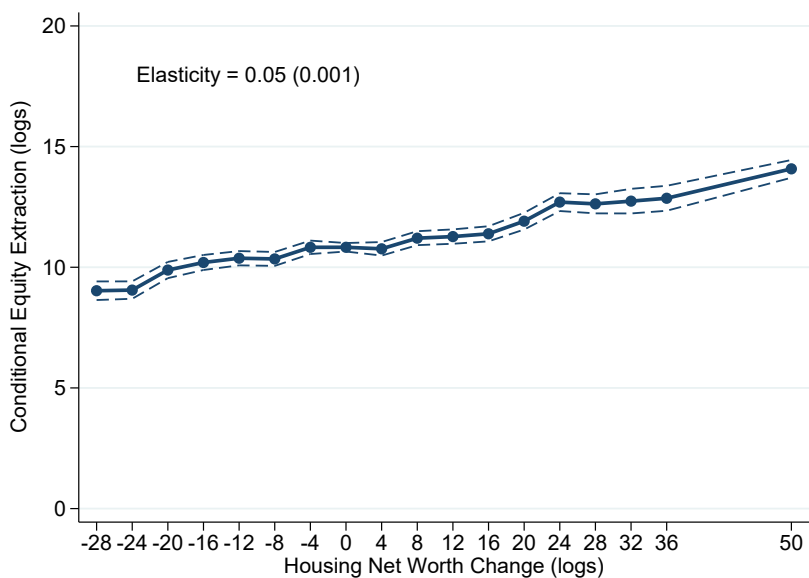
Notes: Panel A plots distributions of residualized house price growth, with and without fixed effects for the last contract duration choice (time until reset) interacted with month and county dummies. The panel shows that past duration choices can explain a large part of the residual price variation (having already absorbed fixed effect for household, month, and county x year). Panel B investigates if past duration choices can also explain residual income variation and shows that it cannot. The fact that past duration is able to predict house price growth, but not other determinants of borrowing such as income, makes it useful for identifying the effects of house prices on borrowing.

FIGURE A.IV: ALTERNATIVE SPECIFICATIONS

A: From Logs to Levels

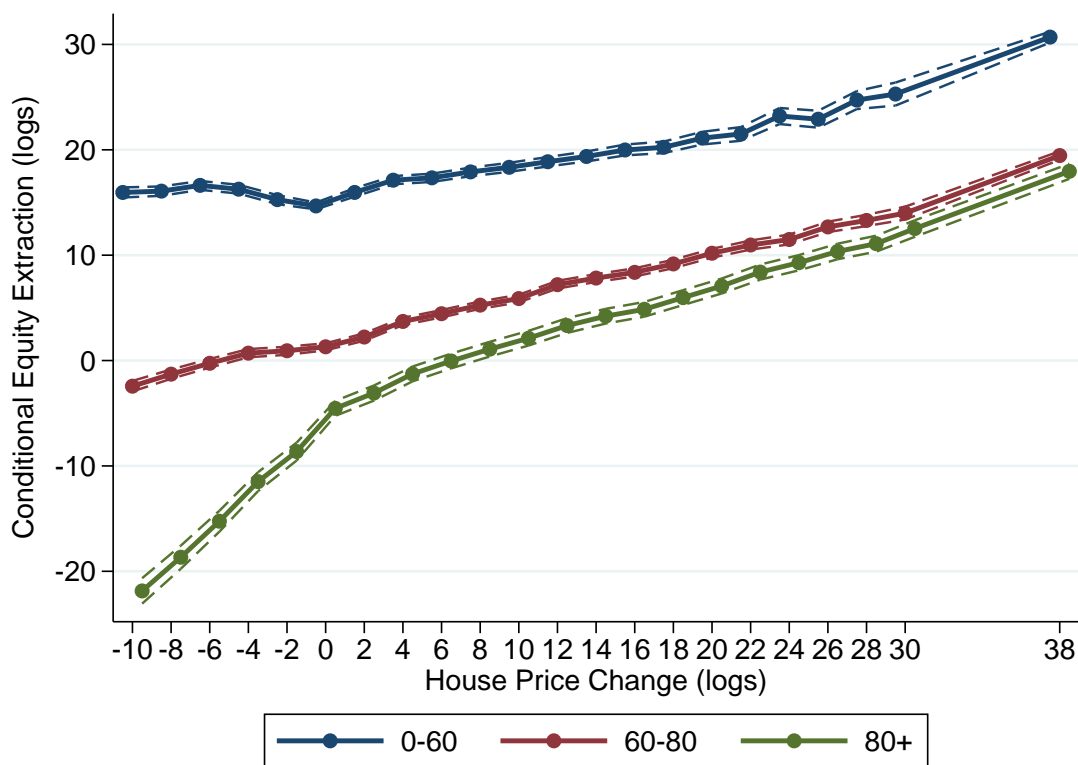


B: From House Prices to Housing Net Worth



Notes: The figure investigates if the previous results are affected by moving from a log-specification to a level-specification (Panel A) and by moving from house prices to housing net worth as the explanatory variable (Panel B). Apart from these changes, the figure is based on the previous fixed effects specification (3) and the panels are constructed in the same way as Figure 10. The results are qualitatively unaffected by the changes, but the alternative specifications are useful for obtaining different parameters. Panel A yields an estimate of the marginal propensity to borrow (equal to 0.11), while Panel B yields an estimate of the equity extraction elasticity with respect to housing net worth (equal to 0.05).

FIGURE A.V: HETEROGENEITY BY LTV NON-PARAMETRICALLY



Notes: The figure plots average equity extraction in different bins of house price growth and in different bins of pre-determined LTV. Pre-LTV is defined as the LTV ratio at time t absent any equity extraction/injection at time t and absent any house price growth between t and $t - 1$. The figure considers three bins of pre-LTV: low leverage (0-60%), intermediate leverage (60-80%), and high leverage (above 80%). The dashed lines represent 95% confidence intervals based on standard errors clustered by household. The figure shows that the *level* of equity extraction decreases with leverage, while the *slope* of equity extraction with respect to house price growth increases with leverage. This is consistent with the collateral channel.

TABLE A.I: EQUITY EXTRACTION ELASTICITIES BY HOME IMPROVEMENT

	(1)	(2)	(3)	(4)
Panel A: Last Mortgage for Home Improvement				
Equity Extraction Elasticity	0.191 (0.006)	0.183 (0.011)	0.171 (0.013)	0.162 (0.014)
Observations	114,566	114,566	108,237	96,613
Panel B: Last Mortgage Not for Home Improvement				
Equity Extraction Elasticity	0.213 (0.003)	0.198 (0.006)	0.189 (0.007)	0.184 (0.007)
Observations	553,200	553,200	524,513	470,038
Panel C: Purpose of Last Mortgage Unknown				
Equity Extraction Elasticity	0.337 (0.002)	0.240 (0.008)	0.250 (0.008)	0.235 (0.009)
Observations	716,580	716,580	678,984	606,975
<u>Control Variables:</u>				
Month FE		×	×	×
Household FE		×	×	×
County x Year FE			×	×
Household Controls				×

Notes: The table reports estimates of the equity extraction elasticity, splitting the estimation sample by whether the last equity extraction decision was made for home improvements or not. Panel A considers homeowners whose last refinance was for home improvements, Panel B considers homeowners whose last refinance was not for home improvements, while panel C considers homeowners whose last refinance purpose is missing in the data. Standard errors are clustered by household and shown in parentheses. The household controls included in column (4) are income level, income growth, the last mortgage interest rate, the age of the borrower, a dummy for couples, and dummies for the various reasons for both the last and current refinance (pure refinance / home improvement / debt consolidation / other). The table shows that, across the different fixed effects specifications, the estimated elasticity is quite stable across samples.