

The Effect of House Prices on Household Borrowing: A New Approach[†]

By JAMES CLOYNE, KILIAN HUBER, ETHAN ILZETZKI, AND HENRIK KLEVEN*

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 the 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. Second, the effect of house prices on borrowing can be explained largely by collateral effects. We study the collateral channel through a multivariate and nonparametric heterogeneity analysis of proxies for collateral and wealth effects. (JEL D14, E32, R31)

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 online Appendix Figure A.I, which shows the evolution of house price growth, consumption growth, and mortgage debt growth in the United States and the United Kingdom 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 (for example, Mian and Sufi 2011, 2014; Mian, Rao, and Sufi 2013; Kaplan, Mitman, and Violante 2017). In this paper we revisit

*Cloyne: Department of Economics, University of California Davis, Davis, CA 95616 (email: jcloyne@ucdavis.edu); Huber: Becker Friedman Institute, University of Chicago, Chicago, IL 60657 (email: kilianhuber@uchicago.edu); Ilzetzki: Department of Economics, London School of Economics, London WC2A 2AE (email: e.ilzetzki@lse.ac.uk); Kleven: Department of Economics, Princeton University, Princeton, NJ 08544 (email: kleven@princeton.edu). Thomas Lemieux was the coeditor for this article. We thank Orazio Attanasio, Richard Blundell, Christopher Carroll, Eric French, Adam Guren, Erik Hurst, Amir Kermani, Thomas Lemieux, Erzo Luttmer, Atif Mian, Magne Mogstad, Emi Nakamura, Ricardo Reis, David Romer, Sebastian Sieglöch, David Sraer, Jón Steinsson, Amir Sufi, Joseph Vavra, Garry Young, three anonymous referees, 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 Regulatory Authority. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <https://doi.org/10.1257/aer.20180086> to visit the article page for additional materials and author disclosure statements.

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 online Appendix Figure A.I does not have a causal interpretation, a point emphasized by Campbell and Cocco (2007); Attanasio et al. (2009); and Attanasio, Leicester, and Wakefield (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 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 United Kingdom 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 United Kingdom, mortgage appraisals provide unbiased measures of actual house prices. Second, the data has a *panel dimension* as many homeowners refinance several times during the 11-year window we consider. This results from the fact that refinancing is a frequent phenomenon in the United Kingdom, because long-term fixed interest mortgages are not available (see Best et al. forthcoming). The panel dimension of the data 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 United Kingdom 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 predetermined, idiosyncratic contract choices as a quasi-experiment for house prices.

¹Much 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, for example, Davidoff 2013, 2016).

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 and Stafford 2004).

We present two 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 (see, e.g., Mian and Sufi 2011), they are also indicative of the potential consumption effects of house prices and they relate to the same underlying mechanisms. We find clear evidence that house price appreciation induces homeowners to increase borrowing by extracting equity from their home. The elasticity of borrowing with respect to house prices lies between 0.2–0.3 and is robust across a range of specifications. We use both fixed effects and instrumental variables (IV) regressions. In our preferred specifications, 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 (predetermined, quasi-exogenous) refinance timing interacts with the housing cycle. Unlike previous studies, many of our results are based on nonparametric, graphical analyses in which we do not impose any a priori assumptions on functional form. A new finding from this approach is that the borrowing elasticity is constant across the distribution of house price changes.³

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, for example, Sinai and Souleles 2005, Berger et al. 2018).⁴ All else equal, the wealth effect should be larger for older 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 and Cocco 2007; Attanasio et al. 2009; Attanasio, Leicester, and Wakefield 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.

We resolve this issue through a multivariate and nonparametric analysis of heterogeneity in the elasticity of borrowing with respect to house prices. We consider four dimensions simultaneously: loan-to-value (LTV), age, income, and income growth. Our approach shows how the borrowing elasticity varies across bins of a given dimension, while simultaneously allowing for differences in the elasticity across bins of the other three dimensions. 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 strong. More levered households are more responsive to house prices, with borrowing elasticities around 0.6 at loan-to-value ratios above 85 percent. By contrast, the age profile is completely flat after controlling nonparametrically for the

³The finding of an isoelastic relationship motivates our focus on log-log specifications through most of the paper, because the log-log coefficient is a direct estimate of the elasticity (in robustness checks, we also report estimates of the marginal propensity to borrow). Reporting the elasticity also eases comparisons to the part of the literature that estimates the elasticity of total borrowing, as opposed to only mortgage borrowing, because there is no mechanical reason why these elasticities should differ. A possible economic reason for the elasticities to differ is that mortgage debt is generally cheaper than other forms of consumer debt, in which case households may shift debt onto their mortgage following a house price increase. Such shifting would lead our elasticity of mortgage borrowing to *overestimate* the elasticity of total borrowing.

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.

other dimensions. The strong relationship between borrowing elasticities and LTV is consistent with evidence on subprime borrowing in the United States (Mian and Sufi 2009, 2011), and it indicates that the collateral channel is the main mechanism behind house price effects.

The UK mortgage market offers an additional way of investigating the collateral channel, arising from the presence of *observable* credit constraints that depend on collateral. Specifically, the UK mortgage interest rate schedule features numerous discrete jumps (notches) at critical LTV thresholds.⁵ We argue that these notches are “soft” collateral constraints, because they represent discrete increases in the cost of borrowing due to a lack of collateral (i.e., due to a high LTV ratio). The only difference between soft borrowing constraints and the hard borrowing constraints familiar from theoretical models is the size of the notch: a hard borrowing constraint is one where the borrowing cost jumps to infinity at a threshold.

For some households, house price growth raised their collateral sufficiently to move them past a lower notch, and thereby reduced their cost of borrowing. For other households, the same change in their house price did not move them past a lower notch, because their initial LTV was located further from a notch. Hence, the UK setting allows us to identify exactly those households, for whom house price growth raised their available collateral in a way that relaxed their cost of borrowing. We find that the borrowing elasticity depends critically on whether the underlying price variation relaxed collateral constraints (by pulling homeowners down to lower notches), reinforced collateral constraints (by pushing homeowners up to higher notches), or left collateral constraints unchanged. In particular, the elasticity is high (around 0.5) among homeowners whose collateral constraint was relaxed by house price growth, and it is zero among those whose collateral constraint was reinforced. Taken together, the heterogeneity analyses using LTV and notches provide evidence that collateral-based changes in the cost of credit play an important role in driving the borrowing response to house price growth.

Given that much of the recent literature focuses on the United States, it is natural to ask if our results are transportable to the US setting. Three points are worth highlighting. First, 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, Mitman, and Violante 2017) are not relevant for assessing external validity in our setting. Second, the majority of the US literature uses cross-regional variation in house prices. Regional effects may differ from our micro elasticities due to local general equilibrium effects, which may amplify or moderate the responses of individual households. Third, institutional differences between the United States and the United Kingdom may lead to differences in the true elasticity of borrowing with respect to house price growth. For example, the elasticity may differ because the fixed costs of equity extraction are higher in the United States. Importantly, however, our empirical approach allows us to accurately capture the entire household mortgage borrowing response to house price changes in the United Kingdom.

⁵Best et al. (forthcoming) describe and analyze these notches in the United Kingdom, while DeFusco and Paciorek (2017) investigate a notch in the US mortgage interest rate schedule.

The paper is organized as follows. Section I reviews the related literature, Section II describes the institutional setting and data, Section III analyzes the sources of house price variation used for identification, Section IV presents results on the effect of house prices on borrowing, Section V presents results on heterogeneity and mechanisms, and Section VI concludes.

I. Literature Review

Important contributions by Mian and Sufi (2011) and Mian, Rao, and Sufi (2013) have shaped the recent debate about the effect of house prices on household debt and consumption. Their findings suggest that house price booms and busts were key determinants of US economic growth before and during the Great Recession. To estimate the effect of house prices, Mian and Sufi (2011) and Mian, Rao, and Sufi (2013) rely on regional house price variation and use housing supply constraints due to topography (from Saiz 2010) to build an instrumental variables (IV) strategy. Kaplan, Mitman, and Violante (2017); Aladangady (2017); and Stroebel and Vavra (forthcoming) use similar IV strategies to study the impact of house prices. Other papers that use regional variation include Campbell and Cocco (2007); Attanasio et al. (2009); Disney, Gathergood, and Henley (2010); Gan (2010); Case, Quigley, and Shiller (2013); and Bhutta and Keys (2016). Studies by Muellbauer and Murphy (1990) and Carroll, Otsuka, and Slacalek (2011), on the other hand, rely on pure time-series variation to estimate the effect of house price growth on borrowing and consumption. Finally, studies by Bostic, Gabriel, and Painter (2009); Disney, Bridges, and Gathergood (2010); Disney and Gathergood (2011); and Cooper (2013) use individual, self-reported house price assessments to estimate the effect of house prices on borrowing and consumption.

Table 1 summarizes existing estimates of how house prices affect borrowing and consumption. The estimates fall in a relatively wide range. A challenge to interpreting the existing results is the possible bias from confounding shocks that are correlated with house price variation across time, regions, and individuals. A number of papers highlight this identification challenge. Attanasio, Leicester, and Wakefield (2011) argue that macroeconomic shocks and expectations explain the correlation between house prices and borrowing. Hurst and Stafford (2004) show that the timing of refinancing is endogenous to household liquidity shocks. Agarwal (2007) finds that households who overestimate their house price are more likely to extract equity and to default on loans. Davidoff (2013, 2016) points out that topography-based instruments are based on a strong exclusion restriction in the US context, because they largely capture variation between the coasts (such as San Francisco, Los Angeles, and New York City) and the interior (such as Wichita, Dayton, and Tulsa).

Motivated by these concerns, we develop an approach that relies on idiosyncratic variation in the timing of refinance events driven by predetermined mortgage contract durations. Our use of contract durations to form an IV strategy is similar in spirit to the pioneering study by Card (1990), who used variation in the duration of Canadian union wage contracts to identify unexpected changes in real wages. Also related, Di Maggio et al. (2017) analyze mortgage contracts that adjust the interest rate after a predetermined duration, but their research question (the impact of interest rates rather than house prices) and empirical strategy is different from ours.

TABLE 1—OVERVIEW OF THE LITERATURE ESTIMATING THE EFFECT OF HOUSE PRICES ON BORROWING AND CONSUMPTION

Citation	Country of study	Outcome	Main explanatory variable	Identifying variation	Elasticity estimate	Level estimate
Aladangady (2017)	USA	Household consumption growth	House price growth	County housing supply elasticity		0.05
Attanasio et al. (2009)	UK	Household consumption growth	House price growth	Regional house price growth	0.04–0.29	
Bhutta and Keys (2016)	USA	Household mortgage borrowing growth	House price growth	Zip-code house price growth		0.07
Bostic, Gabriel, and Painter (2009)	USA	Household consumption	House price	Individual, self-reported house prices	0.04–0.07	
Campbell and Cocco (2007)	UK	Household non-durable consumption growth	House price growth	Regional house price growth	1.22	0.08
Carroll, Otsuka, and Slacalek (2011)	USA	Aggregate consumption growth	House price growth	Aggregate house price growth		0.02–0.09
Case, Quigley, and Shiller (2013)	USA	Consumption per capita growth in the state	Growth in housing market wealth	State house price growth	0.03–0.18	
Cooper (2013)	USA	Household consumption growth	Housing equity growth	Individual, self-reported house price growth		0.06
DeFusco (2018)	USA	Household borrowing growth	Growth in the collateralizable value of houses	Expiry of resale price caps on some houses in a Maryland county		0.04–0.13
Disney and Gathergood (2011)	USA	Household borrowing growth	House price growth	Individual, self-reported house price growth/state house price growth		0.06–0.10
Disney, Bridges, and Gathergood (2010)	UK	Household borrowing growth	House price growth	Individual, self-reported house price growth		–0.01 (0.31 for LTV > 80%)
Disney, Gathergood, and Henley (2010)	UK	Household consumption growth	House price growth	County house price growth, conditional on household expectations		0.01
Gan (2010)	Hong Kong	Growth in household credit card spending	House price growth	District house price growth	0.17	0.02
Kaplan, Mitman, and Violante (2016)	USA	Non-durable expenditures growth	Change in net worth due to house price growth	County housing supply elasticity	0.24–0.36	
Leth-Petersen (2010)	Denmark	Household spending growth	Growth in the collateralizable value of houses	Reform that allowed housing equity to be used as collateral for consumption loans		0.03
Mian and Sufi (2011)	USA	Household borrowing growth	House price growth	County housing supply elasticity	0.52	0.25
Mian, Rao, and Sufi (2013)	USA	Growth in household credit card and auto spending	Change in net worth due to house price growth	County housing supply elasticity	0.6–0.8	0.05–0.07

Notes: The table lists papers that explicitly report estimates for either the elasticity of borrowing/consumption with respect to house prices (in the column *Elasticity estimate*) or the marginal propensity to borrow/consume out of house price changes (in the column *Levels estimate*).

An additional contribution of our paper is that we make no a priori assumptions about the functional form between house prices and borrowing. Using our rich data and nonparametric graphs, we show that the relationship is roughly isoelastic. Our nonparametric and multivariate heterogeneity analysis is also new to the literature and informs an unresolved debate. Previous studies have found a negative age profile of wealth effects, which is inconsistent with standard life-cycle models (Attanasio and

Weber 1994; Attanasio et al. 2009; Attanasio, Leicester, and Wakefield 2011; Mian and Sufi 2011; Bhutta and Keys 2016; Berger et al. 2018). We show that the negative age profile reflects the confounding effects of collateral and that the true age profile is flat.

Finally, our paper highlights the importance of collateral constraints in driving the effect of house prices on borrowing. A number of the papers listed in Table 1 report that more leveraged households respond to house price growth more strongly. We add to these findings with our multivariate heterogeneity analysis, which shows the effects of leverage are not driven by other dimensions of heterogeneity correlated with leverage. A related literature analyzes the effect of relaxed access to housing collateral on borrowing (Leth-Petersen 2010; DeFusco 2018), retail sales (Abdallah and Lastrapes 2012), consumption (Agarwal and Qian 2017), and entrepreneurship (Jensen, Leth-Petersen, and Nanda 2014).⁶ Compared to these studies, our approach allows us to examine not only the effects of a collateral shock, but more generally how households respond to a house price shock, including tests for the wealth channel. In addition, we use a large and representative sample (the population of UK mortgagors), study the effects of both relaxing and tightening collateral constraints, and introduce the analysis of notches as new test of the collateral channel. Complementing our findings, two recent papers argue that LTV-dependent borrowing constraints affect households in response to other shocks, such as debt reductions (Ganong and Noel 2018) and changes in mortgage payments (Di Maggio et al. 2017).

II. Institutional Setting and Data

A. 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 United Kingdom. Almost all mortgage products feature a relatively low interest rate for an initial period, followed 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 much lower rate. For example, at current rates a refinancer 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.

In addition to the penalizing reset after the end of the initial low-interest period, most mortgage contracts feature large early repayment charges, typically 5 or

⁶The identifying variation in these studies comes from government policy. For example, DeFusco (2018) analyzes the expiration of resale price caps on houses in a county in Maryland, which increased the housing collateral available to homeowners. His estimates for the marginal propensity to borrow out of housing collateral lie between 0.04 and 0.13.

⁷More than 90 percent of mortgage products feature such reset rate structures (see, for example, MoneyFacts.co.uk).

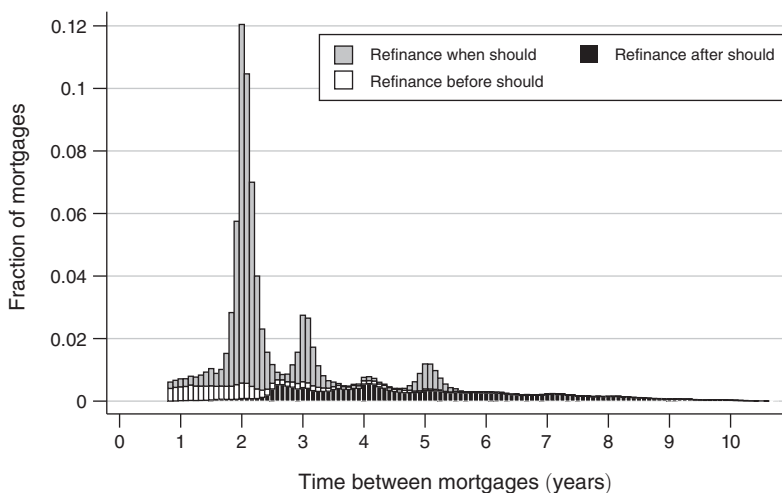


FIGURE 1. HOMEOWNERS REFINANCE AROUND THE ONSET OF THE RESET RATE

Notes: The figure shows the distribution of the time between mortgage financing events. Households who refinance between two months before and six months after the onset of their reset rate are shown in light gray, households who refinance more than six months after the onset of their reset rate are shown in black, and households who refinance more than two 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.

10 percent of the outstanding loan. These charges make it very costly to refinance or adjust borrowing before the end of the initial period.

The combination of penalizing reset rates and heavy early repayment charges implies that households have strong incentives to refinance right around the end of the initial duration. To confirm that households act on these incentives, Figure 1 shows the distribution of time between mortgages among refinancers 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 two months before and six 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 and Stafford 2004). Second, the fact that refinance events are frequent allows

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.

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 United Kingdom. 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 percent of all household debt in the United Kingdom. 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 percent, 70 percent, 75 percent, 80 percent, and 85 percent. Online Appendix Figure A.II shows the average interest rate schedule as a function of LTV across all mortgage products (see Best et al. forthcoming 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" collateral 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 homeowner'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.

B. 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 United States (for example, Ben-David 2011; Agarwal, Ben-David, and Yao 2015; Agarwal, Ambrose, and Yao 2017), 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 United Kingdom, as we demonstrate in two ways. First, while we do not observe actual market prices for refinanced properties, we do observe market prices (along with

⁹Best et al. (forthcoming) provide a bunching analysis of borrowing responses to these interest 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 percent.

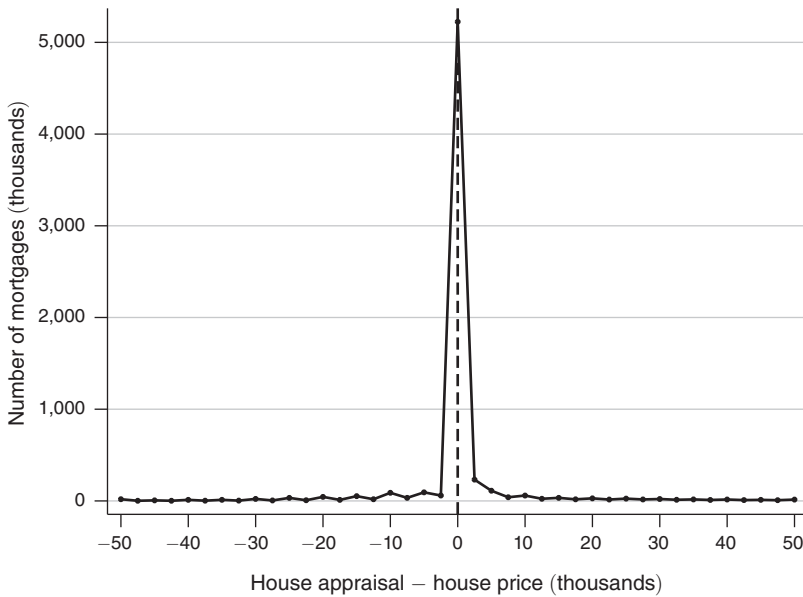


FIGURE 2. HOUSE PRICES VERSUS APPRAISALS (NEW PURCHASES)

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.

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 appraised prices are 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 nonparametrically for two observables: the age of the homeowner and the postcode of the property. Specifically, we run the following regression separately for the purchase and appraisal price series:

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




FIGURE 3. HOUSE PRICES VERSUS 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 20 quantiles of the age distribution as well as 20 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.

where the first term includes a full set of quarter dummies, the second term includes dummies for 20 quantiles of the age distribution, and the third term includes dummies for 20 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, and it includes 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 nonparametric 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.¹¹

C. Data

The data come from a new and comprehensive regulatory dataset containing the universe of mortgage product sales. These data are collected by the UK 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

¹¹Further evidence against consequential appraisal bias is that the equity extraction elasticity remains stable when controlling for fixed effects for month, household, and county \times year, as well as a number of time-varying household characteristics (results in Section VA). Typical sources of appraisal bias are that certain households or banks tend to demand biased appraisals or that region- or household-specific income shocks lead to biased appraisals. The control variables account for all these possibilities.

completed household mortgage product originations from April 2005, but does not include commercial or buy-to-let mortgages.¹²

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 age, gross income, and whether the income is solely or jointly earned.¹³

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 one-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 estimation sample is therefore a subset of the refiners in the PSD, for which we have at least two mortgage observations. Some of our specifications below control for individual fixed effects, so they identify solely off refiners with at least three mortgage observations (for which we can calculate the house price change and equity extraction for at least two refinancing events).

Table 2 summarizes the data. Panel A compares descriptive statistics for home buyers (column 1), all refiners (column 2), refiners in our estimation sample with at least two mortgage observations in the PSD (column 3), and refiners in our estimation sample with at least three mortgage observations in the PSD (column 4). There are no significant differences between any of the groups in the share of couples, income, income growth, interest rate, and house price. Some differences between buyers and refiners are to be expected. For example, buyers tend to be younger and have higher LTV ratios.

Panel B of Table 2 reports statistics for the 1.38 million observations in our estimation sample with at least two observations, split into three subsamples. As discussed above, practically all mortgages in the United Kingdom 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

¹² See <https://www.fca.org.uk/firms/product-sales-data> for officially published, high-level data.

¹³ Full details of the dataset can be found on the FCA's PSD website.

TABLE 2—DESCRIPTIVE STATISTICS

	Buyers (1)	Refinancers (2)	<i>Refinancers in our estimation sample</i>	
			With ≥ 2 observations (3)	With ≥ 3 observations (4)
<i>Panel A. Buyers versus refinancers</i>				
Age	36.47 (10.13)	42.08 (9.77)	40.85 (8.90)	41.40 (8.06)
Couple	0.52 (0.50)	0.54 (0.50)	0.54 (0.50)	0.56 (0.50)
Income	55,282.17 (556,583.42)	55,949.83 (145,816.42)	57,602.96 (81,440.65)	58,412.07 (53,227.65)
Income change (logs)		0.08 (0.36)	0.08 (0.35)	0.06 (0.32)
Interest rate	4.39 (1.40)	4.51 (1.40)	3.98 (1.50)	3.84 (1.39)
House price	229,375.32 (326,209.46)	248,328.76 (361,735.65)	256,517.10 (187,020.25)	266,479.30 (181,871.70)
LTV	70.72 (21.67)	56.53 (21.80)	61.50 (18.96)	59.61 (17.53)
Observations	7,119,807	5,935,441	1,384,346	305,232
	Refinance on-time	Refinance off-time	Missing duration	
<i>Panel B. Refinancers in our estimation sample with ≥ 2 observations</i>				
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	54,516.32 (48,424.02)	53,442.66 (52,355.65)	62,005.61 (108,733.95)	
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	245,030.89 (163,127.94)	233,110.00 (158,358.87)	276,638.16 (213,289.69)	
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 statistics for home buyers (column 1), all refinancers (column 2), refinancers in our estimation sample with at least two observations (column 3), and refinancers in our estimation sample with at least three observations (column 4). Panel B compares statistics for three subsamples of our estimation sample with at least two observations: 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 onset of the reset rate.

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.

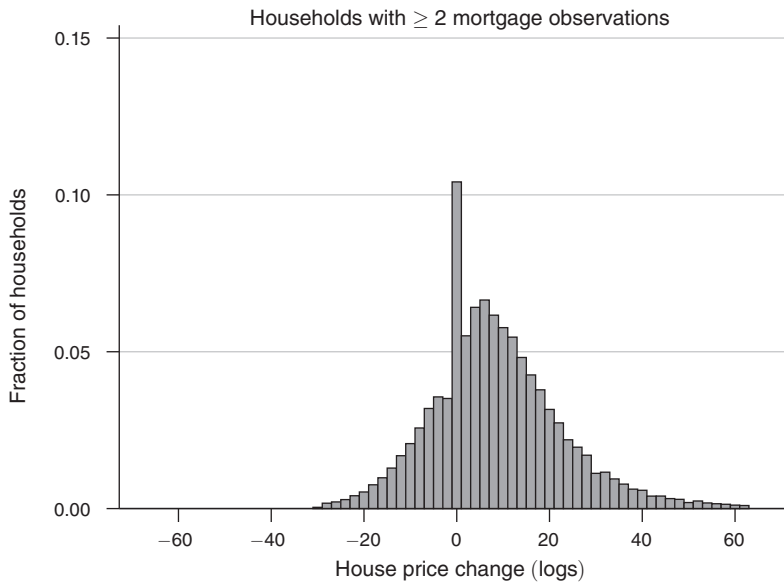


FIGURE 4. DISTRIBUTION OF RAW HOUSE PRICE GROWTH

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).

III. 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 percent and $+60$ percent 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 United Kingdom 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

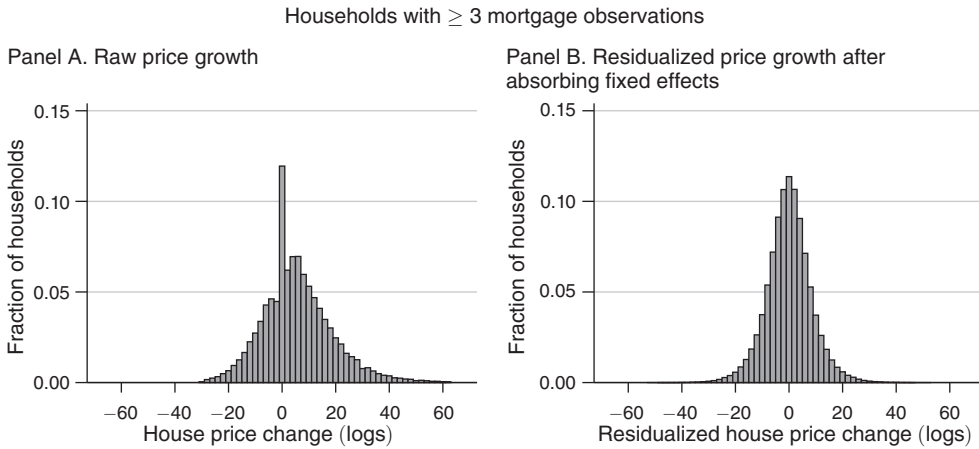


FIGURE 5. DISTRIBUTION OF RAW VERSUS RESIDUALIZED HOUSE PRICE GROWTH

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).

growth gives an R^2 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 percent and $+20$ percent 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 allows 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

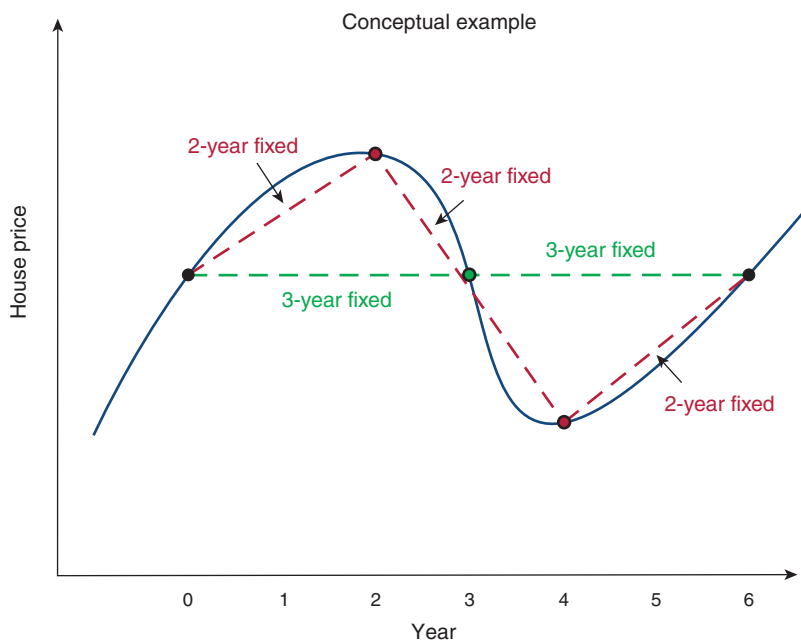


FIGURE 6. THE TIMING OF REFINANCE EVENTS AND HOUSE PRICE CHANGES

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 two-year fixed interest rate loans, while the other prefers three-year fixed interest loans. The homeowner in two-year contracts refinances three times over a six-year period, facing either positive or negative price shocks at each event, whereas the homeowner in three-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.

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

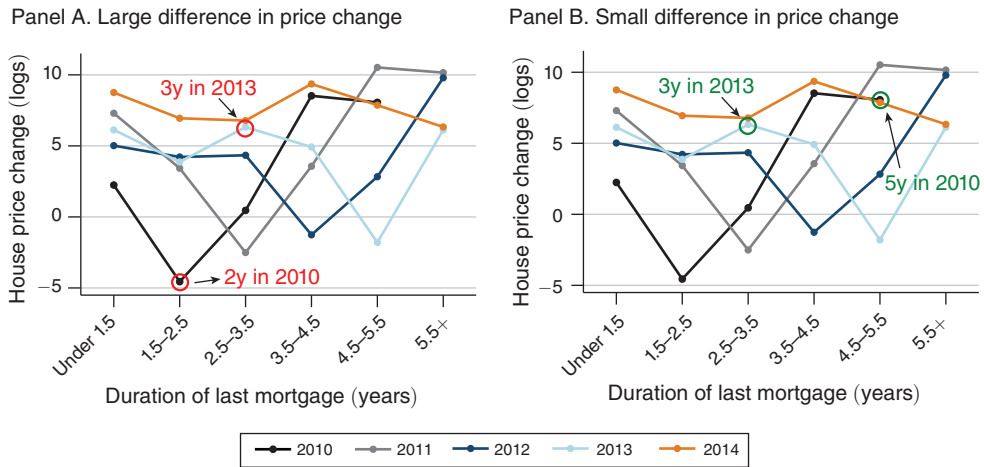


FIGURE 7. HOUSE PRICE CHANGES VERSUS LAST DURATION \times TIME OF REFINANCE

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.

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 (for example, Case and Shiller 1989; Shiller 2007; Case, Shiller, and Thompson 2012; Kaplan, Mitman, and Violante 2017). 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, online Appendix Figure A.III 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 estimate the borrowing response to house price growth using two types of strategies. We first consider ordinary least squares (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.

IV. Do House Prices Affect Borrowing?

A. OLS Fixed Effects Specification

The outcome variable in our analysis is the amount of equity extracted at the time of refinancing. We define this outcome as the log difference between the debt a homeowner holds after refinancing and the debt a homeowner would have held had she simply rolled over the preexisting debt when refinancing, i.e., without extracting or injecting any equity. This outcome is given by $\log D_{ict} - \log D_{ict}^P$, where D_{ict} denotes mortgage debt of individual i in county c at refinance time t and D_{ict}^P denotes the predetermined debt at time t based on past debt choices and amortization.¹⁴

To investigate the effect of house price growth on equity extraction, we specify

$$(2) \quad \log D_{ict} - \log D_{ict}^P = \sum_j \beta_j \cdot \mathbf{1}[\Delta \log P_{it} \in j] + \alpha_i + \gamma_t + \delta_{ct} + X_{it} \theta + \nu_{ict},$$

where P_{it} denotes the price of the house owned by individual i at time t . Note that we consider a nonparametric specification in which 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. While we primarily consider log-specifications, we will also explore level-specifications and show that those yield the same qualitative results.¹⁵ We allow for individual fixed effects α_i , time fixed effects γ_t (at the monthly

¹⁴That is, we have $D_{ict}^P = D_{ict-1} + (\text{amortization between } t-1 \text{ and } t)$.

¹⁵The 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 translated into an

level), and county-by-time fixed effects δ_{ct} (at the yearly level).¹⁶ The county-by-time fixed effect absorbs regional, time-varying factors (such 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 specification has the form of a triple-differences specification relying on within-individual variation in price growth. Note that X_{it} includes a number of individual, time-varying variables that could be relevant for debt demand.

We begin the analysis by plotting the estimated coefficients $\hat{\beta}_j$ in different bins of house price growth, leaving out the other controls in equation (2). Of course, this is equivalent to plotting the raw averages of equity extraction across the different bins of house price growth. The results are shown in panel A of Figure 8. Three insights are worth highlighting. First, overall there is a clear positive relationship between house price growth and equity extraction. We see that equity extraction increases from 5–10 percent of debt to almost 25 percent of debt as house price growth changes from –10 percent to +40 percent. 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 0.23. This elasticity masks the heterogeneity between the negative and positive ranges of house price growth, with an elasticity of 0.4 in the positive range.

While the raw patterns in panel A are consistent with an impact of house prices on borrowing, the relationship may be affected by the confounding effects on borrowing that we have discussed. For example, the asymmetry between negative and positive house price growth could reflect such confounders. Therefore, panel B of Figure 8 considers the results from a richer specification that controls for individual fixed effects, time fixed effects, and county-by-time fixed effects. Interestingly, 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.¹⁷ The average borrowing elasticity is 0.2, slightly lower than the previous estimate.¹⁸

These findings are robust to alternative specifications, which we demonstrate through a number of checks presented in the online Appendix. Figure A.IV

average borrowing elasticity in the population in order to compare with the log-specification).

¹⁶Counties correspond to UK local planning authorities (as described above). There is some abuse of notation in specification (2) as we use t to describe time in both months and years.

¹⁷The asymmetry disappears because this graph plots house price growth conditional on individual, time, and county-by-time fixed effects. The asymmetry still exists when we look at negative house price growth in absolute terms. For example, in a specification with all the fixed effects, the equity extraction elasticity among households with absolute house price gains is 0.28 (0.01), while the elasticity among households with absolute house price declines is 0.01 (0.04). These findings are consistent with the view that liquidity constraints prevent homeowners from injecting equity when house prices fall.

¹⁸Note that all of our estimates include both extensive margin effects (whether or not to extract equity) and intensive margin effects (how much equity to extract, conditional on extracting). The results in online Appendix Table A.I show that our estimates are driven primarily by the intensive margin. There is only a very small extensive margin effect of house price growth on the probability of (strictly) positive equity extraction.

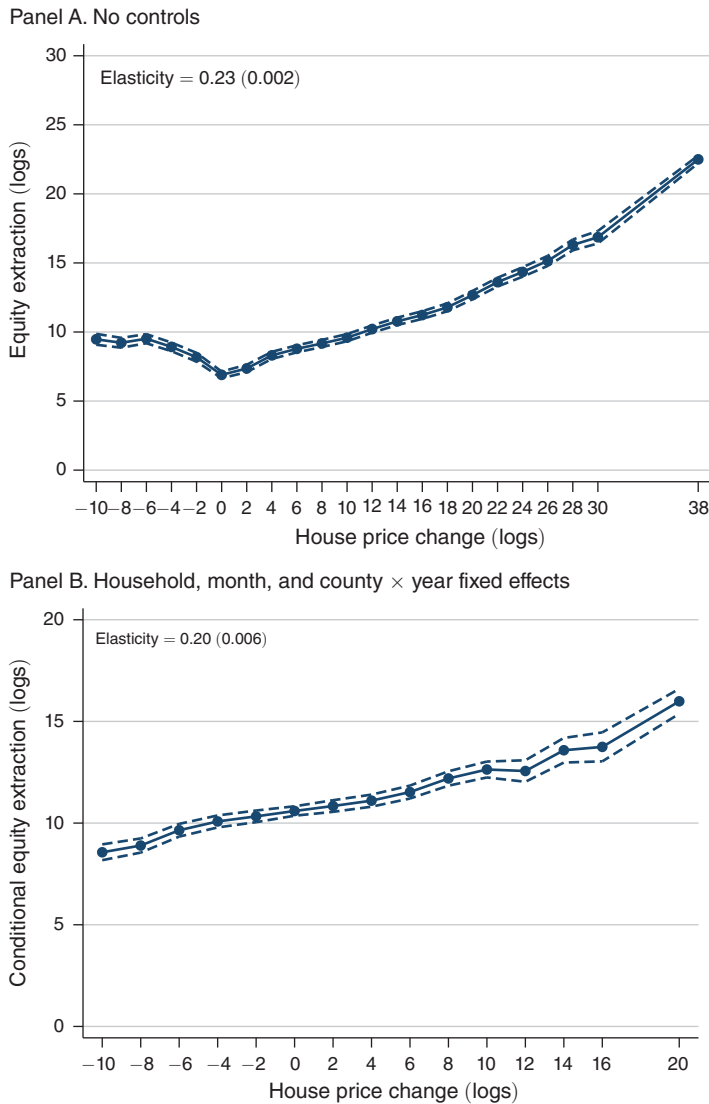


FIGURE 8. EQUITY EXTRACTION VERSUS HOUSE PRICE GROWTH

Notes: Panel A plots the average equity extraction in different bins of house price growth. 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 it simply rolled over the existing mortgage debt at the time of refinancing), multiplied by 100. House price growth is measured as log change between refinance events, multiplied by 100. Panel B plots conditional equity extraction in different bins of house price growth based on the fixed effects specification (2). 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. The other covariates are fixed effects for household, month, and county \times year. The dashed lines represent 95 percent confidence intervals based on standard errors clustered by household. Each panel reports the average equity extraction elasticity based on a log-linear specification. Panel B shows an almost perfectly log-linear relationship between equity extraction and house prices, conditional on the covariates.

shows that the relationship between equity extraction and house price changes remains log-linear and similarly sloped both in a more parsimonious specification (dropping county-by-time fixed effects) and in a richer specification (adding

time-varying, household controls).¹⁹ While the borrowing elasticity is not affected by these specification changes, it does feature a modest degree of cyclicity as we show in online Appendix Figure A.V. The largest elasticities are observed in the run-up to the recession and the smallest elasticities are observed in the middle of the recession (see also Guren et al. 2018 for an analysis of time variation in the effects of house price changes).²⁰

As discussed in Sections II and III, our empirical strategy is based on the idea that the timing of refinance is quasi-exogenous in the United Kingdom. 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 IIA 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 3 presents estimates of borrowing elasticities across samples that vary by refinance timing: the full sample in panel A (summarizing the results already presented in figures), 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 United Kingdom 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.²¹ Overall, the table shows that elasticity estimates are 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 somewhat higher in the off-time sample, consistent with a small selection bias.

Given that the existing literature has relied on regional variation in house prices, it is interesting to investigate the implications of using spatial variation in our context. For this exercise, we use annual house price growth at the county level as our treatment variable, and we use annualized equity extraction as our outcome variable (i.e., equity extraction normalized by the number of years between the current and previous refinance event). As we show in online Appendix Table A.II, the borrowing elasticities are much larger when estimated from regional variation than those

¹⁹The household-level controls included in panel B of the figure are income level, income growth, the last mortgage interest rate, age of the borrower, a dummy for couples, and dummies for a range of self-reported reasons for the current and the last refinance.

²⁰We have also investigated alternative specifications for the outcome and treatment variables. Starting from equation (2), online Appendix Figure A.VI shows how the 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). Panel A yields an estimate of the marginal propensity to borrow (equal to 0.11) and panel B yields an estimate of the elasticity with respect to housing net worth (equal to 0.05).

²¹To 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.

TABLE 3—EQUITY EXTRACTION ELASTICITIES BY REFINANCE TIMING

	(1)	(2)	(3)	(4)
<i>Panel A. Full sample</i>				
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
<i>Panel B. On-time sample</i>				
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
<i>Panel C. Off-time sample</i>				
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
<i>Panel D. Sample with missing durations</i>				
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
<i>Control variables</i>				
Month fixed effects		×	×	×
Household fixed effects		×	×	×
County × year fixed effects			×	×
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 two months before and six months after reset rate onset), panel C considers the sample of off-time refinancers (defined as those who refinance more than two months before or more than six 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).

obtained from our preferred specifications above. This is likely due to the fact that geographic variation in house prices is correlated with various confounders, and we have no compelling instrument for geographic house price variation in the UK context.²²

Finally, we investigate the issue of home improvements. Variation in individual house prices may be partly driven by idiosyncratic home improvements, which are endogenous and may not represent true increases in household net worth.²³

²²Hilber and Vermeulen (2016) construct a topography-based housing supply elasticity index for England (à la Saiz 2010), but not for the rest of the United Kingdom (Northern Ireland, Scotland, and Wales). However, besides the issues with the exclusion restriction of such instruments (as discussed above), Hilber and 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.

²³In particular, home improvements do not increase household net worth unless they increase the house price by more than the amount invested in the house.

As a first check, we use self-reported information on the reason for refinancing that is available for part of the sample. Online Appendix Table A.III 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 (panel C). The table shows that the estimated elasticity is quite stable across samples. Specifically, among those who report no home improvement, the elasticities are similar to the elasticities for the full sample discussed above. This alleviates any major concerns about home improvements, but we acknowledge that our measure of home improvements is imperfect. Hence, the next section goes further by presenting IV estimates that cannot be plausibly affected by home improvements.

B. 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, for example due to 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 (predetermined) 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, the most compelling source of variation is a three-way interaction between the duration choice in the last mortgage (say 2-year versus 3-year fixed interest rate), the time of the current refinance event (say 2010 versus 2011), and the regional house price cycle. Hence, we construct instruments based on the interaction between dummies for past duration choices, dummies for the time of refinance, and dummies for different regions. The first stage of the IV is specified as

$$\begin{aligned}
 (3) \quad \Delta \log P_{it} = & \textit{duration}_{it} \otimes \textit{region}_i \otimes \textit{year}_t + \textit{duration}_{it} + \alpha_i + \gamma_t \\
 & + \textit{duration}_{it} \otimes \textit{region}_i + \textit{duration}_{it} \otimes \textit{year}_t \\
 & + \textit{region}_i \otimes \textit{year}_t + X_{it} \eta + \mu_{it},
 \end{aligned}$$

where \otimes denotes the outer product, so that the instrumental variables ($\textit{duration}_{it} \otimes \textit{year}_t \otimes \textit{region}_i$) 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). X_{it} is a vector of individual, time-varying control variables.

TABLE 4—EQUITY EXTRACTION ELASTICITIES USING INSTRUMENTAL VARIABLES

	(1)	(2)	(3)	(4)
IV equity extraction elasticity	0.150 (0.004)	0.284 (0.026)	0.268 (0.078)	0.249 (0.076)
Observations	769,494	769,494	769,494	768,042
<i>Control variables</i>				
Month fixed effects		×	×	×
Household fixed effects		×	×	×
Contract duration fixed effects		×	×	×
Contract duration × region fixed effects			×	×
Contract duration × year fixed effects			×	×
Region × year fixed effects			×	×
Household controls				×

Notes: The table reports estimates of the equity extraction elasticity using instrumental variables (IV). The instruments are interactions of dummies for the last mortgage contract duration (time until reset), year, and region. The table shows IV elasticities from four different specifications, with the richest specification in column 4 corresponding to equations (3)–(4). 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). Standard errors are clustered by household and given in parentheses. The IV elasticities are close to the elasticities from the OLS fixed effects specifications.

The second stage of the IV is

$$\begin{aligned}
 (4) \quad \log D_{it} - \log D_{it}^P &= \beta \cdot \widehat{\Delta \log P_{it}} + duration_{it} + \alpha_i + \gamma_t \\
 &+ duration_{it} \otimes region_i + duration_{it} \otimes year_t \\
 &+ region_i \otimes year_t + X_{it}\eta + \nu_{it},
 \end{aligned}$$

where $\widehat{\Delta \log P_{it}}$ is the predicted house price growth from the first-stage specification (3).

Our richest specification includes fixed effects for household, month, and duration dummies on their own.²⁴ It also includes fixed effects for all the second-level interactions of the dummies that we use to construct the instrumental variables. These second-level interactions are fixed effects for duration-by-region, duration-by-year, and region-by-year. Because we include these control variables, the identification in the richest specification comes entirely from the three-way interaction between last duration, region, and year. This overcomes a number of key identification concerns. The last duration dummies on their own eliminate the possibility that differences in risk preferences across households drive the effects. The duration-by-region dummies ensure that differences across regions in the types of households that choose certain durations do not influence the results. The duration-by-year dummies eliminate the possibility that time-varying shocks to households with certain contract duration bias the estimates.

²⁴Household and month dummies subsume region and year dummies, respectively.

Table 4 shows the estimated elasticities of equity extraction with respect to house price across four IV specifications. The richest specification in column 4 corresponds to the specifications shown in equations (3)–(4). There is a nontrivial difference in the estimates between the basic specification without any controls in column 1 and the richer specifications in columns 2–4. But across the richer specifications, the IV elasticity estimates are stable (around 0.25–0.28) and marginally 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 results presented so far: there is a clear positive effect of house prices on borrowing, with a borrowing elasticity between 0.2–0.3.

V. 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. (2018) provide a theoretical foundation for the various mechanisms that may be at play. 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 and Cocco 2007; Case, Quigley, and Shiller 2013). However, it is not obvious that such changes in nominal wealth translate into real wealth, as highlighted by Sinai and 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 life-cycle. 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 and Weber 1994, Campbell and Cocco 2007, Attanasio et al. 2009, Mian and 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 (for example, Aoki, Proudman, and Vlieghe 2004; Iacoviello 2005), and it has been argued to be empirically important for household borrowing in a number of studies (for example, Lustig and van Nieuwerburgh 2005, Mian and Sufi 2011, DeFusco 2018).²⁵ The collateral channel

²⁵The collateral channel has also been shown to be important for business investments and employment (Chaney, Sraer, and Thesmar 2012; Adelino, Schoar, and Severino 2015).

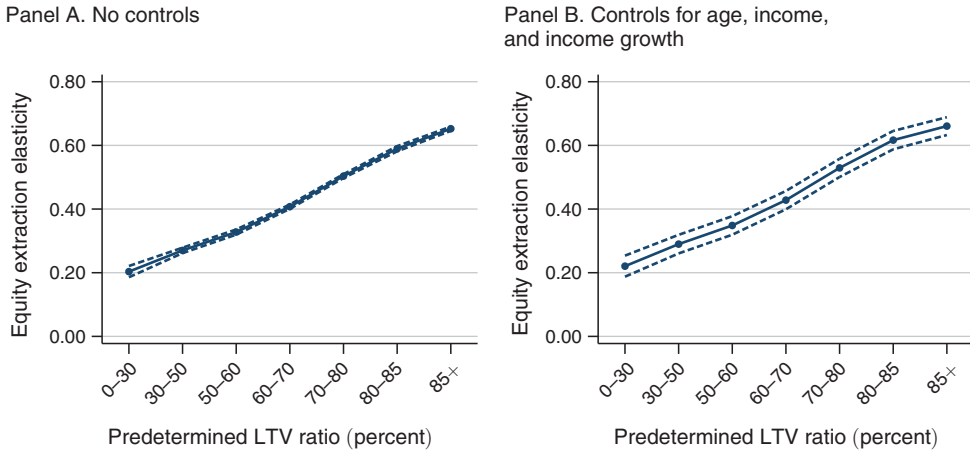


FIGURE 9. HETEROGENEITY IN BORROWING ELASTICITY BY LTV

Notes: The figure shows heterogeneity in the equity extraction elasticity by LTV. 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 $t - 1$. Panel A is based on a univariate specification that investigates heterogeneity by LTV on its own, while Panel B is based on a multivariate specification allowing for heterogeneity in four dimensions simultaneously: LTV, age, income level, and income growth. The multivariate specification is shown in equation (5). The dashed lines give 95 percent confidence intervals based on standard errors clustered by household. The figure shows a strong, increasing relationship between LTV and the borrowing elasticity, consistent with collateral effects.

implies heterogeneity across leverage ratios: households 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 start by disentangling wealth and collateral effects based on a heterogeneity analysis that uses the power and granularity of our administrative data. We conduct a multivariate and nonparametric analysis of heterogeneity in the borrowing elasticity along the main dimensions predicted to determine household borrowing responses. 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 different method to assess its empirical importance.

A. Heterogeneity Analysis

We investigate how the borrowing elasticity varies along four dimensions of heterogeneity: loan-to-value (LTV), 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

$$(5) \log D_{it} - \log D_{it}^P = \sum_k \sum_j \beta_j^k \cdot \mathbf{1}[X_{it}^k \in j] \cdot \Delta \log P_{it} + \sum_k \sum_j \lambda_j^k \cdot \mathbf{1}[X_{it}^k \in j] + \nu_{it}$$

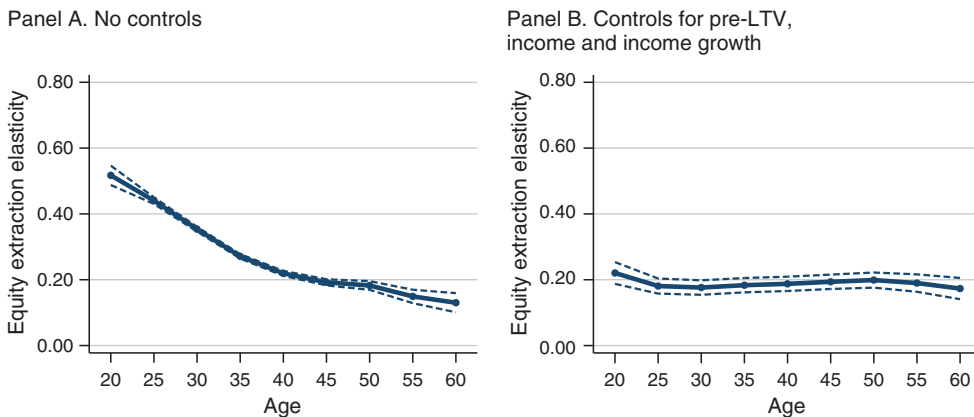


FIGURE 10. HETEROGENEITY IN BORROWING ELASTICITY BY AGE

Notes: The figure shows heterogeneity in the equity extraction elasticity by age. Panel A is based on a univariate specification that investigates heterogeneity by age on its own, while Panel B is based on a multivariate specification allowing for heterogeneity in four dimensions simultaneously: LTV, age, income level, and income growth. The multivariate specification is shown in equation (5). The dashed lines give 95 percent confidence intervals based on standard errors clustered by household. The figure shows a negative or flat relationship between age and the borrowing elasticity, inconsistent with wealth effects.

where $\mathbf{1}[X_{it}^k \in j]$ is a dummy equal to 1 when variable k (LTV, age, income, or income growth) falls in bin j . By allowing for a large set of bin dummies in each dimension (seven LTV bins, nine age bins, seven income bins, and seven income growth bins), and by allowing for these dummies to affect both the slope and the intercept, our analysis is very nonparametric. 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 (5) 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 similar, but standard errors increase substantially in fixed effects specifications with heterogeneity.

Figure 9 investigates heterogeneity with respect to LTV, which is the main proxy for collateral effects, as discussed above. We analyze heterogeneity by predetermined 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 ratio 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 (5), which we show in a fully nonparametric specification in online Appendix Figure A.VII.²⁶

Figure 10 studies the effect of age. These graphs show heterogeneity in the borrowing elasticity across 5-year bins between the ages of 20 and 60. Panel A presents results without controls for the other dimensions of heterogeneity. The figure shows the opposite pattern than what is suggested by standard life-cycle 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. 2018 for a similar argument). Panel B 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, online Appendix Figures A.VIII and A.IX display heterogeneity across income levels and income growth, respectively. 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 in panel A and the multivariate specification in panel B. These graphs 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 nonparametrically for other correlated factors, 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 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 (5). A 1 percent increase in the house price represents a 5 percent increase in housing net worth for a homeowner at 80 percent LTV, but only a 2 percent increase for a homeowner at 50 percent LTV. Mechanically, there are heterogeneous wealth changes depending on LTV. As a robustness check, we have therefore tried a level specification as well, finding similar qualitative results. This strengthens the conclusion that the collateral channel is crucial. Third and last, leverage may be correlated with unobserved individual characteristics that affect borrowing behavior. A candidate would be self-control problems. As Mian and

²⁶Online Appendix Figure A.VII presents nonparametric 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 percent), intermediate-leverage homeowners (LTV between 60–80 percent), and high-leverage homeowners (LTV above 80 percent). 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 life-cycle. 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.

Sufi (2011) note, it is likely that households with greater self-control problems will be observed as credit constrained. However, when augmenting equation (5) 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).

To conclude, the heterogeneity results suggest that the collateral channel plays a central role in driving the borrowing elasticity. In the next section, we propose a new test, one that exploits discrete changes in the tightness of collateral constraints around interest notches, providing another piece of evidence in favor of the collateral channel.

B. Collateral Channel: A Test Using Interest Notches

The UK setting offers a new way of testing the importance of collateral as driver of the borrowing elasticity. This test is motivated by the insight that a *hard* collateral constraint can be defined as a sharp increase to infinity (or a prohibitive level) in the cost of borrowing when collateral falls below some threshold (or similarly, when LTV surpasses a threshold). Following this definition, any sharp increase in the cost of borrowing that depends on the amount of collateral is also a collateral constraint. We view any collateral-based threshold where the cost of borrowing increases discontinuously, without reaching infinity, a *soft* collateral constraint.

The UK mortgage interest rate schedule features a number of such soft collateral constraints. These soft collateral constraints are discrete increases in the interest rate (interest notches) at critical LTV thresholds, as described in Section IIIA. There are interest notches at LTV ratios of 50 percent, 60 percent, 70 percent, 75 percent, 80 percent, 85 percent, and 90 percent.²⁷ As the LTV ratio surpasses (and housing collateral therefore falls below) one of the critical thresholds, the cost of borrowing increases sharply. The only difference between such interest notches and the hard borrowing constraints from theoretical models is the size of the notch: a hard borrowing constraint is one where the borrowing cost jumps to infinity at a threshold.²⁸

The direct incentive created by these interest notches is for homeowners to choose LTV ratios just below one of the thresholds, thus creating bunching in the LTV distribution. Such bunching represents borrowing responses to the interest rate, rather than responses to the house price, and was studied by Best et al. (forthcoming) for the United Kingdom and DeFusco and Paciorek (2017) for the United States. The focus of our exercise here is different. We consider what happens when house price growth, by increasing the available collateral, moves homeowners above or below interest notches and sharply changes the cost of borrowing. If collateral constraints play a limited role in driving the borrowing response to house price growth, households should extract equity to a similar degree independent of whether house price growth moves their LTV above or below an interest notch. But if collateral constraints play an important role, equity extraction should differ depending on whether households' LTV is moved above or below an interest notch.

²⁷ Online Appendix Figure A.II illustrates most of these notches.

²⁸ In fact, the notch at 90 percent LTV 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.

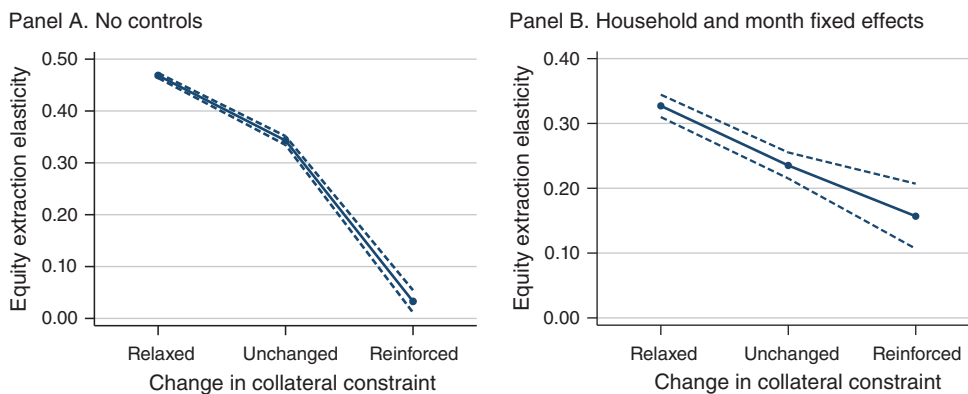


FIGURE 11. HETEROGENEITY IN BORROWING ELASTICITY BY NOTCHES MOVED

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 percent, 60 percent, 70 percent, 75 percent, 80 percent, 85 percent, and 90 percent. 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 percent confidence intervals based on standard errors clustered by household.

Figure 11 investigates how the equity extraction elasticity depends on whether the underlying price variation moves homeowners across notches. 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.”²⁹ Panel A considers a baseline specification without any other controls. This is a specification like equation (5), 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.³⁰ 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.

One concern is that the pattern of heterogeneity in panel A of Figure 11 could be due to the fact that households experiencing negative house price growth do not have a positive elasticity to house price growth in a specification without controls (as exhibited in panel A of Figure 8). We therefore carry out a heterogeneity test using a specification that does not exhibit any underlying difference between households experiencing positive versus negative house price changes. Panel A of online

²⁹ However, this terminology should not be taken literally: house price appreciation may relax credit constraints even if it does not move homeowners to a lower interest notch.

³⁰ While the figure pools all years 2005–2015, we have checked that the patterns are roughly the same inside and outside the recession years.

Appendix Figure IV shows that, conditional on household and month fixed effects, households experiencing positive or negative house price growth respond similarly. Panel B of Figure 11 therefore introduces household and month fixed effects in the specification. This graph confirms the qualitative relationship between the borrowing 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 soft collateral constraints underscores the importance of the collateral channel. In online Appendix B, we develop yet another test of the collateral channel. There 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. We show that, when house price growth pulls homeowners below an interest notch and relaxes their credit constraint, they respond by increasing their LTV up to the notch. This interaction between house price growth and bunching responses is again consistent with collateral-based borrowing decisions.

VI. Conclusion

The global financial crisis has reignited a debate about the role of house prices in driving household borrowing. A first generation of papers following the crisis studied this question using regional data from the United States and found strong borrowing responses. This paper takes a different methodological tack to the question. Using administrative mortgage data and unique features of the UK mortgage market, we develop an empirical approach that relies on idiosyncratic variation in the timing of refinance events driven by predetermined mortgage contract durations. We find that a 10 percent increase in individual house prices raises borrowing by around 2–3 percent.

Our rich dataset also allows us to explore why borrowing responds to house prices. The striking finding from a heterogeneity 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 nonparametrically for factors such as age, income, and income growth. This heterogeneity analysis, together with a test using soft collateral constraints (interest rate notches based on collateral), suggests that the housing collateral channel is the main driver of the elasticities we find.

The magnitude of these responses, and the importance of collateral constraints, has important implications for understanding household behavior in both micro- and macroeconomics. 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

- Abdallah, Chadi S., and William D. Lastrapes. 2012. "Home Equity Lending and Retail Spending: Evidence from a Natural Experiment in Texas." *American Economic Journal: Macroeconomics* 4 (4): 94–125.

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino.** 2015. "House Prices, Collateral, and Self-Employment." *Journal of Financial Economics* 117 (2): 288–306.
- Agarwal, Sumit.** 2007. "The Impact of Homeowners' Housing Wealth Misestimation on Consumption and Saving Decisions." *Real Estate Economics* 35 (2): 135–54.
- Agarwal, Sumit, Brent W. Ambrose, and Vincent W. Yao.** 2017. "The Effects and Limits of Regulation: Appraisal Bias in the Mortgage Market." Unpublished.
- Agarwal, Sumit, Itzhak Ben-David, and Vincent Yao.** 2015. "Collateral Valuation and Borrower Financial Constraints: Evidence from the Residential Real Estate Market." *Management Science* 61 (9): 2220–40.
- Agarwal, Sumit, and Wenlan Qian.** 2017. "Access to Home Equity and Consumption: Evidence from a Policy Experiment." *Review of Economics and Statistics* 99 (1): 40–52.
- Aladangady, Aditya.** 2017. "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." *American Economic Review* 107 (11): 3415–46.
- Aoki, Kosuke, James Proudman, and Gertjan Vlieghe.** 2004. "House Prices, Consumption, and Monetary Policy: A Financial Accelerator Approach." *Journal of Financial Intermediation* 13 (4): 414–35.
- Attanasio, Orazio P., Laura Blow, Robert Hamilton, and Andrew Leicester.** 2009. "Booms and Busts: Consumption, House Prices and Expectations." *Economica* 76 (301): 20–50.
- Attanasio, Orazio P., Andrew Leicester, and Matthew Wakefield.** 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.
- Attanasio, Orazio P., and Guglielmo Weber.** 1994. "The UK Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence." *Economic Journal* 104 (427): 1269–302.
- Ben-David, Itzhak.** 2011. "Financial Constraints and Inflated Home Prices during the Real Estate Boom." *American Economic Journal: Applied Economics* 3 (3): 55–87.
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra.** 2018. "House Prices and Consumer Spending." *Review of Economic Studies* 85 (3): 1502–42.
- Best, Michael, James Cloyne, Ethan Ilzetzki, and Henrik Kleven.** Forthcoming. "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches." *Review of Economic Studies*.
- Bhutta, Neil, and Benjamin J. Keys.** 2016. "Interest Rates and Equity Extraction during the Housing Boom." *American Economic Review* 106 (7): 1742–74.
- Bostic, Raphael, Stuart Gabriel, and Gary Painter.** 2009. "Housing Wealth, Financial Wealth, and Consumption: New Evidence from Micro Data." *Regional Science and Urban Economics* 39 (1): 79–89.
- Campbell, John Y., and João F. Cocco.** 2007. "How Do House Prices Affect Consumption? Evidence from Micro Data." *Journal of Monetary Economics* 54 (3): 591–621.
- Card, David.** 1990. "Unexpected Inflation, Real Wages, and Employment Determination in Union Contracts." *American Economic Review* 80 (4): 669–88.
- Carroll, Christopher D., Misuzu Otsuka, and Jiri Slacalek.** 2011. "How Large Are Housing and Financial Wealth Effects? A New Approach." *Journal of Money, Credit, and Banking* 43 (1): 55–79.
- Case, Karl E., John M. Quigley, and Robert J. Shiller.** 2013. "Wealth Effects Revisited: 1975–2012." *Critical Finance Review* 2: 101–28.
- Case, Karl E., and Robert J. Shiller.** 1989. "The Efficiency of the Market for Single-Family Homes." *American Economic Review* 79 (1): 125–37.
- Case, Karl E., Robert J. Shiller, and Anne K. Thompson.** 2012. "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets." *Brookings Papers on Economic Activity* 43 (2): 265–98.
- Chaney, Thomas, David Sraer, and David Thesmar.** 2012. "The Collateral Channel: How Real Estate Shocks Affect Corporate Investment." *American Economic Review* 102 (6): 2381–409.
- Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven.** 2019. "The Effect of House Prices on Household Borrowing: A New Approach: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20180086>.
- Cooper, Daniel.** 2013. "House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral." *Review of Economics and Statistics* 95 (4): 1183–97.
- Davidoff, Thomas.** 2013. "Supply Elasticity and the Housing Cycle of the 2000s." *Real Estate Economics* 41 (4): 793–813.
- Davidoff, Thomas.** 2016. "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated with Many Demand Factors." *Critical Finance Review* 5 (2): 177–206.
- DeFusco, Anthony A.** 2018. "Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls." *Journal of Finance* 73 (2): 523–73.

- DeFusco, Anthony A., and Andrew Paciorek.** 2017. "The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit." *American Economic Journal: Economic Policy* 9 (1): 210–40.
- Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, and Amit Seru.** 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107 (11): 3550–88.
- Disney, Richard, Sarah Bridges, and John Gathergood.** 2010. "House Price Shocks and Household Indebtedness in the United Kingdom." *Economica* 77 (307): 472–96.
- Disney, Richard, and John Gathergood.** 2011. "House Price Growth, Collateral Constraints and the Accumulation of Homeowner Debt in the United States." *B.E. Journal of Macroeconomics*: 11 (1): 1–30.
- Disney, Richard, John Gathergood, and Andrew Henley.** 2010. "House Price Shocks, Negative Equity, and Household Consumption in the United Kingdom." *Journal of the European Economic Association*: 8 (6): 1179–207.
- Gan, Jie.** 2010. "Housing Wealth and Consumption Growth: Evidence from a Large Panel of Households." *Review of Financial Studies* 23 (6): 2229–67.
- Ganong, Peter, and Pascal Noel.** 2018. "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." NBER Working Paper 24964.
- Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson.** 2018. "Housing Wealth Effects: The Long View." NBER Working Paper 24729.
- Hilber, Christian A. L., and Wouter Vermeulen.** 2016. "The Impact of Supply Constraints on House Prices in England." *Economic Journal* 126 (591): 358–405.
- Hurst, Erik, and Frank Stafford.** 2004. "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption." *Journal of Money, Credit, and Banking* 36 (6): 985–1014.
- Iacoviello, Matteo.** 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* 95 (3): 739–64.
- Jensen, Thais Lærkholm, Søren Leth-Petersen and Ramana Nanda.** 2014. "Housing Collateral, Credit Constraints and Entrepreneurship: Evidence from a Mortgage Reform." NBER Working Paper 20583.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante.** 2016. "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data." NBER Working Paper 22232.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante.** 2017. "The Housing Boom and Bust: Model Meets Evidence." NBER Working Paper 23694.
- Kleven, Henrik.** 2016. "Bunching." *Annual Review of Economics* 8: 435–64.
- Leth-Petersen, Søren.** 2010. "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?" *American Economic Review* 100 (3): 1080–103.
- Lustig, Hanno N., and Stijn G. van Nieuwerburgh.** 2005. "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective." *Journal of Finance* 60 (3): 1167–219.
- Mian, Atif, Kamalesh Rao, and Amir Sufi.** 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–726.
- Mian, Atif, and Amir Sufi.** 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." *Quarterly Journal of Economics* 124 (4): 1449–96.
- Mian, Atif, and Amir Sufi.** 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review* 101 (5): 2132–56.
- Mian, Atif, and Amir Sufi.** 2014. *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. Chicago: University of Chicago Press.
- Muellbauer, John, and Anthony Murphy.** 1990. "Is the UK Balance of Payments Sustainable?" *Economic Policy: A European Forum* 5 (11): 347–95.
- Saiz, Albert.** 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.
- Shiller, Robert J.** 2007. "Understanding Recent Trends in House Prices and Homeownership." In *Federal Reserve Bank of Kansas City Economic Policy Symposium Proceedings*, 89–123.
- Sinai, Todd, and Nicholas S. Souleles.** 2005. "Owner-Occupied Housing as a Hedge against Rent Risk." *Quarterly Journal of Economics* 120 (2): 763–89.
- Stroebel, Johannes, and Joseph Vavra.** Forthcoming. "House Prices, Local Demand, and Retail Prices." *Journal of Political Economy*.