Interest Rates, Debt and Intertemporal Allocation: Evidence From Notched Mortgage Contracts in the UK*

Michael Carlos Best, Stanford University
James Cloyne, Bank of England
Ethan Ilzetzki, London School of Economics
Henrik Jacobsen Kleven, London School of Economics

July 2015

Abstract

Using a novel source of quasi-experimental variation in interest rates, we study the response of household debt and intertemporal consumption allocation to interest rates. We also develop a new approach to structurally estimate the Elasticity of Intertemporal Substitution (EIS). In the UK, the mortgage interest rate schedule features discrete jumps—notches—at thresholds for the loan-to-value (LTV) ratio, creating strong incentives for bunching below those thresholds. We document large and sharp bunching below every notch, which translates into sizable interest elasticities of mortgage debt, between 0.1 and 1.4 across different LTV levels. We develop a dynamic model that links these reduced-form responses to the underlying structural EIS. The EIS is much smaller and less heterogeneous than the reduced-form elasticities, between 0.05-0.25 across LTV levels and household types. We show that our structural approach is robust to a wide range of assumptions on beliefs about the future, uncertainty, risk aversion, discount factors and present bias. Our findings have implications for the numerous calibration studies in economics that rely on larger values of the EIS.

*Best: mbest@stanford.edu, Cloyne: James.Cloyne@bankofengland.co.uk, Ilzetzki: E.Ilzetzki@lse.ac.uk, Kleven: H.J.Kleven@lse.ac.uk. We thank Raj Chetty, Jeff Clemens, Gordon Dahl, Rebecca Diamond, Roger Gordon, Bob Hall, Ethan Kaplan, David Laibson, Emi Nakamura, Petra Persson, Ricardo Reis, Emmanuel Saez, Orie Shelef, Jón Steinsson, Garry Young, Stephen Zeldes, and numerous seminar participants for helpful comments and discussion. 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. All charts and estimates use data provided by the Financial Conduct Authority and MoneyFacts.

1 Introduction

We study the response of household debt and intertemporal consumption allocation to interest rates, and develop a new approach to estimating the Elasticity of Intertemporal Substitution (EIS). Our specific focus is on the demand for mortgage debt, which accounts for the bulk of household debt in most countries. The parameters we estimate are central to a large range of policy questions in household finance, macro and public economics, but there exists no conclusive evidence on their magnitudes due to the difficulties of finding exogenous variation in interest rates (e.g. Gruber 2006).

Our starting point is a novel source of quasi-experimental variation in interest rates arising from the fact that UK banks offer *notched* mortgage interest schedules. That is, the mortgage interest rate features discrete jumps (notches) at specific thresholds for the loan-to-value ratio (LTV), namely at 60%, 70%, 75%, 80%, and 85%. For example, the interest rate increases by almost 0.4%-points at the 80% LTV threshold, meaning that the homeowner pays an additional 0.4%-points annually on the *entire* mortgage when crossing this threshold. These interest rate jumps create strong incentives for households to reduce leverage to a LTV below the notch, thereby giving up consumption today in order to get a lower interest rate and hence more consumption in the future. Such behavioral responses lead to bunching below notches and missing mass above notches in the LTV distribution.

To study these interest rate notches, we use administrative mortgage data from the Financial Conduct Authority. The data cover the universe of household mortgages in the UK between 2008-2014, including rich information on mortgage contracts and borrower characteristics. The majority of UK mortgage products carry a relatively low interest rate for a period of 2-5 years after which a much higher reset rate kicks in, creating very strong incentives to remortgage at the time the reset rate starts to apply. This makes remortgaging a common occurrence. We focus on the leverage choices of remortgagors as this population allows for a clean assessment of debt demand and intertemporal consumption choices. Specifically, because remortgagors have already made their housing choices and their house values are assessed by the lender, estimating LTV responses in this sample allows us to isolate debt choices from housing choices.

We find large bunching below notches and missing mass (holes) above notches in the LTV dis-

tribution. To precisely estimate the amount of bunching and missing mass, we need an estimate of the counterfactual LTV distribution—what the distribution would have looked like without notches—and the public finance literature has developed approaches and discussed problems of obtaining such counterfactuals (Chetty *et al.* 2011; Kleven & Waseem 2013). We develop a new approach to obtain the counterfactual distribution that exploits the panel structure of the remortgagor data. Our approach uses the previous mortgage, its amortization schedule, and the lender's house price appraisal to predict each remortgagor's LTV absent responses to interest rate notches. We show that this yields a compelling non-parametric estimate of the *global* counterfactual distribution. This contrasts with the local estimations in the previous bunching literature and allows us to asses how interest rate notches affect the entire distribution of household leverage.

Based on the counterfactual LTV distribution and the amount of bunching, we can estimate reduced-form elasticities of mortgage demand with respect to the marginal interest rate. Mortgage demand elasticities are sizable, around 0.3 on average, and increasing in leverage. The most highly leveraged households have elasticities larger than 1. These elasticities are central for evaluating a range of policies—such as monetary policy or tax policy—that affect the after-tax interest rate in the mortgage market. A general issue with such reduced-form elasticities, however, is that they depend on many factors of the environment and therefore may not represent any structural parameter or have much external validity. This motivates our structural analysis of the elasticity of intertemporal substitution (EIS).

We develop an intertemporal model of consumption and mortgage demand that allows us to characterize the relationship between the reduced-form responses to interest rate notches and the underlying structural EIS. We estimate that the EIS is about 0.1 on average and that it ranges between 0.05 and 0.25 across LTV levels and household types.[1] Hence, there is only modest heterogeneity in the EIS, with larger elasticities for more leveraged households, older households, households with larger loan-to-income ratios, and those that saw larger house price appreciation. Overall, our estimates are at the lower end of the wide range estimated in the previous literature, most of which have relied on potentially endogenous time series movements in interest rates.

While the structural analysis requires us to make parametric assumptions, our approach is in fact very robust to those assumptions due to the nature of the identifying interest rate variation. We show that our estimates barely move when varying the discount factor, future interest rates,

[1]The finding that the structural EIS is much lower and less heterogeneous than the reduced-form debt elasticity echoes insights from classic calibration studies, showing that a given value of the EIS can imply widely different, but typically much larger, savings elasticities depending on other calibrated parameters (Summers 1981; Evans 1983).

uncertainty and risk aversion (using an Epstein-Zin-Weil formulation that separates the EIS from risk aversion), or when allowing for present-biased borrowers. Roughly speaking, the reason why these factors have no first-order impact is that they mainly create level effects on lifetime utility that are present on both sides of the notch and therefore approximately cancel out in the decision to bunch or not.

Why do we find relatively low structural elasticities? Is it because the true underlying responsiveness is low, or is it because households are affected by optimization frictions (inattention, misperception, etc.) or other constraints that attenuate our estimates? Three points suggest the former. First, our approach explicitly adjusts for optimization frictions using an extension of the frictions methodology developed by Kleven & Waseem (2013). This approach exploits observed density mass in a dominated region above notches to quantify the amount of friction. Second, while liquidity constraints may not be fully captured by our friction adjustment, their impact is likely to be limited in our setting due to the fact that most people extract sizable amounts of housing equity when refinancing. As we show, it would be feasible for most remortgagors to bunch at the notch by taking out less equity. Third, because remortgaging is a common occurrence in the UK, our estimates are not based on a selected sample of households that face unusual constraints or have unusual preferences. For these reasons, we interpret the low EIS estimates as capturing what they are supposed to, namely that intertemporal preferences over consumption are not very price elastic.

Our paper relates to several literatures. First, we contribute to the household finance literature that studies the determinants of household debt and leverage. Although mortgages account for the vast majority of household debt and therefore plays a key role in this literature, there exists very little causal evidence on the responsiveness of mortgage debt to the cost of borrowing.[2] Previous attempts have used variation in the after-tax interest rate generated by government policies such as taxes, subsidies, and regulation (e.g. Follain & Dunsky 1997; Martins & Villanueva 2006; Jappelli & Pistaferri 2007; DeFusco & Paciorek 2014). The work by DeFusco & Paciorek (2014) comes closest to our study by estimating leverage responses to a notch created by the conforming loan limit in the US, but their reduced-form estimates do not separate mortgage demand and housing demand effects, nor do they pursue the analysis of structural parameters as we do here.

Second, we contribute to a large literature studying consumption and savings responses to

[2]There is a richer literature on other smaller components of household debt such as credit cards and car loans (e.g. Gross & Souleles 2002; Attanasio *et al.* 2008).

the after-tax interest rate. This includes a reduced-form literature estimating savings responses to interest rates and savings policies (reviewed by Bernheim 2002) and a structural literature estimating elasticities of intertemporal substitution (reviewed by Attanasio & Weber 2010). The latter estimates consumption Euler equations using either aggregate data (e.g. Hall 1988; Campbell & Mankiw 1989) or micro survey data (e.g. Zeldes 1989; Attanasio & Weber 1993, 1995; Vissing-Jørgensen 2002; Gruber 2006). The main conceptual differences between our approach and this literature is that we use interest rate notches at a point in time as opposed to interest changes over time, and that our estimating equation is not a standard Euler equation due to the discontinuous nature of the notched incentive.

Third, we contribute to a recent literature in public economics that has developed bunching approaches to study behavioral responses to government interventions (Saez 2010; Chetty *et al.* 2011; Kleven & Waseem 2013). This literature initially focused on taxes and labor supply in static contexts, but is now spreading to other areas and to dynamic settings. By combining a bunching approach with dynamic structural estimation, our paper is conceptually related to recent work by Einav *et al.* (2015) on US Medicare. Conceptually the notches we study are different from the policy notches that have been analyzed in the existing literature, because they are introduced by the private sector rather than by policy makers, but the sharp discontinuities nevertheless allow us to estimate parameters that are crucial for policy.

Finally, we contribute to an emerging micro-macro literature, which uses micro data to address questions in macroeconomics (e.g. Johnson *et al.* 2006; Agarwal *et al.* 2007; Mian & Sufi 2009, 2011, 2012; Chodorow-Reich 2014; Auclert 2015; Best & Kleven 2015; Cloyne *et al.* 2015; Di Maggio *et al.* 2015; Keys *et al.* 2015). In contrast to this literature, we move beyond reduced-form evidence by structurally estimating the EIS, a key parameter in all macroeconomic models. The relatively small elasticities we estimate have implications for the importance of intertemporal consumption substitution as a propagation mechanism for business cycles and for the effects of macroeconomic stabilization policies (e.g., Summers 1981; King & Rebelo 1990; Mishkin 1995; Auclert 2015).

The paper is organized as follows. Section 2 describes the institutional setting and data, section 3 outlines our empirical methodology, section 4 presents reduced-form evidence, section 5 develops our model and presents structural estimations, and finally section 6 concludes.

2 Institutional Setting and Data

2.1 Institutional Setting: Notches and Refinancing

Several institutional features facilitate our analysis. First and foremost is the menu of mortgage products that banks offer their customers in the UK. These mortgage products have interest rates that vary across brackets of the LTV ratio, creating discrete jumps in the interest rate—notches—at specific LTV thresholds. The size and location of these notches differ across banks and over time. It is very common for banks to impose large notches at LTV ratios of 75%, 80%, and 85%, and many banks also impose notches at 60%, 70%, and 90%.[3] The LTV notches are very salient: daily newspapers display menus of interest rates by bank and LTV ratios on a regular basis, and the LTV thresholds feature very prominently when shopping for mortgages as we show below.

Second, UK mortgage products come with either a variable interest rate or a fixed interest rate for a relatively short period—ranging from two to five years—after which the homeowner has to refinance to avoid a much higher reset rate. The majority of mortgages are of the fixed-rate variant and homeowners typically remortgage around the time that the reset rate kicks in. This is confirmed by Figure A.1, which shows the distribution of time-to-remortgage and distinguishes between those who remortgage when 'they should' (i.e., around the time that the reset rate kicks in) and before or after they should. The vast majority of mortgage holders refinance when the reset rate starts to apply after 2, 3 or 5 years. The 2-year fixed rate product is the most common in the UK and so the largest spike in the distribution appears at 2 years. Hence, remortgaging is a very common occurrence in the UK, and in this paper we will focus specifically on the remortgagors as this will allow us to isolate debt choices from housing choices and to assess intertemporal consumption substitution. Note that the interest rate notches described above affect the rate charged until the next mortgage (i.e., typically 2-5 years after) as opposed to the rate charged over the entire term of the mortgage (typically 25-35 years).

To illustrate the decision environment facing a mortgagor, this link shows an example of the screen that the customer first sees when shopping for mortgages at this bank, and it is immediately clear that the mortgage products depend on the LTV bracket. Specifically, this bank has LTV brackets of 60-70, 70-75, 75-80, 80-85, and 85-90. Clicking through to a specific product ("2 Year Fixed Standard") we see the information that the customer sees when clicking on this specific product

[3] Following the financial crisis, very few banks have offered mortgages above an LTV of 90% (implying that this threshold becomes a corner rather than a notch for much of our data period) and so we have relatively few observations above 90%. Our empirical analysis therefore focuses on the notches below 90%.

under the relevant LTV range in the first screen. The interest rate at time of writing equaled 3.14% below an LTV of 90%, 2.29% below an LTV of 85%, and 1.99% below an LTV of 80%. This means that the interest rate notches for this particular bank, product and point in time (May 2015) are equal to 0.85%-points at the 85% threshold and 0.3%-points at the 80% threshold. As mentioned, the size of these notches vary by bank, product, and time, and the availability of the different products sometimes depend on borrower characteristics (income, age, family size, etc.). Our empirical analysis will be based on the average interest rate jump at each threshold (conditioning on bank, product, time, and borrower characteristics), which we quantify in section 2.3.

Finally, these webpages illustrate two other facts that help explain the time-to-refinance pattern in Figure A.1. First, the much higher reset interest rate after two years is shown in column 4. The reset rate is listed as 3.94% at the time of writing, but this is a variable rate that could be different than the actual reset rate the mortgagor would end up paying. The large and variable reset rate gives a very strong incentive to refinance at the end of the fixed-rate period. Second, column 6 shows that there is a significant mortgage booking fee (£999 for this product), which means that the homeowner does not want to refinance too often. This, together with the presence of penalties for refinancing early (not shown in this screen), explains why very few homeowners in Figure A.1 refinance before they should. These two facts imply that the time of refinancing is effectively imposed on homeowners by the structure of the mortgage market, and this is helpful for ruling out potential selection issues arising from endogenous remortgage timing decisions.

2.2 Data

Our analysis uses a novel and comprehensive regulatory dataset containing the universe of mortgage product sales collected by the UK's Financial Conduct Authority. This Product Sales Database (PSD) has information on all completed household mortgage product originations back to April 2005, but does not include commercial or buy-to-let mortgages.[4] The scale of the dataset allows us to focus attention on particular types of households and mortgage products that are most suited for our purposes while still retaining a large sample.

Regulated lenders are required to submit quarterly information on all mortgage products originated. This includes a wide range of information about the mortgage such as the loan size, the date the mortgage became active, the valuation of the property, the initial interest rate charged, whether the interest rate is fixed or variable, the end date of the initial interest rate (the time at which the

[4]See http://www.fca.org.uk/firms/systems-reporting/product-sales-data/ for officially published high level data.

higher reset rate starts applying), whether the mortgage payments include capital, 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, whether the income is solely or jointly earned, whether the borrower is a first time buyer, mover or re-mortgagor and the reason for the re-mortgage. There are also some characteristics of the property such as the type of dwelling and the number of rooms.[5]

While we observe the borrower's LTV ratio, the PSD does not include information on product origination fees. These fees, while small relative to the loan size, can sometimes be rolled into the loan without affecting the LTV statistic used to determine the borrower's interest rate. It is therefore possible to observe an actual LTV ratio of for example 75.01% where the borrower was still offered the product with a maximum LTV of 75%. In order to address this, we exploit information on all mortgage products (including origination fees) in the UK available from the organization MoneyFacts between 2008Q4 and 2014Q4.[6] For a mortgage observation in the PSD we find the corresponding product in MoneyFacts based on the lender, the date of the loan, the interest rate offered, and the mortgage type. Where the interest rate paid accords with the lower LTV bracket on offer from the lender, but where their actual LTV is marginally above the threshold, we subtract the product fee from the loan value reported in the PSD. Inspecting these individuals, the additional loan amounts often correspond precisely to the value of the fee, and so this adjustment places a large number of such individuals exactly at the notch. While this matching exercise reduces the sample, it is crucial for our methodology that the LTV ratio we use corresponds exactly to the one determining the actual interest rate.

Another useful feature of the PSD is that we are able to observe whether the household is remortgaging. Using information about the characteristics of the property and the borrower, we can match-up remortgagors over time to construct a panel. As described later, the panel structure allows us to implement a novel approach for estimating the counterfactual LTV distribution absent notches. The remortgagor panel will therefore be the baseline data set for our analysis.

Table 1 shows a broad range of descriptive statistics in different samples. Column 1 includes the full sample of mortgages sold between 2008Q4 and 2014Q4 where we can exploit fee information from MoneyFacts. The full sample is large, with around 2.8 million observations. Column 2 shows how the properties of the sample change when we restrict attention to remortgagors. The

[5] Full details of the dataset can be found on the FCA's PSD website.

[6] www.moneyfacts.co.uk.

descriptive statistics are very similar, although the LTV and LTI ratios are slightly lower for remortgagors as one would expect. Column 3 shows the descriptive statistics in the panel of remortgagors that we use in the empirical analysis. In moving from column 2 to 3 we lose remortgagors where we lack sufficient information about their previous loans as well as those we are not able to match-up over time. Our estimation sample still includes over 550,000 mortgages. Importantly, the descriptive statistics are very stable across the three columns, suggesting that our estimation sample (column 3) has similar average characteristics as the full population of mortgagors.

2.3 Interest Rate Jumps at Notches

As described above, the UK mortgage market features discrete interest rate jumps at critical LTV thresholds, namely at 60%, 70%, 75%, 80%, and 85%. The first step of our analysis is to assess the size of these interest rate notches. Unlike standard bunching approaches in which the notched incentive is the same for all individuals (for example a tax rate notch as in Kleven & Waseem 2013), in our setting the interest rate notch varies across households depending on lender, mortgage product, the time of loan origination, and potentially individual characteristics. Our empirical analysis will use the average interest rate jump at each notch conditional on lender, product, time, and individual characteristics. We estimate these interest rate jumps non-parametrically using a very flexible regression.

In our baseline specification, we estimate the interest rate schedule based on the following regression:

$$\begin{aligned} r_i &= f\left(LTV_i\right) + \beta_1 \text{lender}_i + \beta_2 \text{type}_i \otimes \text{dur}_i \otimes \text{month}_i + \beta_3 \text{repayment}_i + \beta_4 \text{reason}_i \\ &\quad + s_1\left(\text{age}_i\right) + s_2\left(\text{income}_i\right) I\left\{\text{single}_i\right\} + s_3\left(\text{income}_i\right) I\left\{\text{couple}_i\right\} + s_4\left(\text{term}_i\right) + \nu_i \quad (1) \end{aligned}$$

where r_i is the nominal mortgage interest rate for individual i, $f\left(.\right)$ is a step function with steps at each 0.25%-point LTV bin, lender_i is a vector of lender dummies, type_i is a vector of mortgage type dummies (fixed interest rate, variable interest rate, capped interest rate, and "other"), dur_i is a vector of dummies for the duration of the initial mortgage period (the time until the reset rate kicks in), month_i is a vector of dummies for the month in which the mortgage was originated, repayment_i is a vector of dummies for the repayment type (interest only, capital and interest, and "other"), reason_i is a vector of dummies for the reason for the remortgage (pure remortgage, home improvement, debt consolidation, and "other"), and $s_j\left(.\right)$, $j = 1, ..., 4$ are cubic splines with knots

at the quintiles of the distributions of age, income (separately for singles and couples), and the total term length (e.g., 30 for a 30-year mortgage). We denote by $\otimes$ the outer product, and so the term $\text{type}_i \otimes \text{dur}_i \otimes \text{month}_i$ allows for each combination of product type and duration to have its own non-parametric time trend.

Figure 1 plots the conditional interest rate as a function of LTV based on specification (1) for the full sample. That is, in each LTV bin we plot the coefficient on the LTV bin dummy plus a constant given by the predicted value $E[\hat{r}_i]$ at the mean of all the other covariates (i.e., omitting the contribution of the LTV bin dummies). The figure shows that the mortgage interest rate evolves as a step function with sharp jumps at LTV ratios of 60%, 70%, 75%, 80%, and 85%. The jumps in the interest rate at the three upper notches are very large and roughly similar in size, around 0.4%-points of extra annual interest above each notch. The jumps in the interest rate at the two lower notches are smaller, especially at the 60% threshold where the jump is about 0.12%-points. Between notches the interest rate schedule is very flat, which is a great advantage for our identification strategy: this implies that, conditional on the variables we observe, the mortgage interest rate is almost fully determined by the LTV notches that we exploit. In fact, the interest rate schedule in Figure 1 looks qualitatively similar to the tax rate schedules that researchers have used for identification in numerous settings, but we are not aware of previous work that have made use of such schedules to study the effects of interest rates.

In Figure 1 we control for individual characteristics such as age, income, and family status, but these have very little effect on the mortgage interest rate conditional on lender and product characteristics. Figure A.2 in the appendix shows the results from a specification that excludes all individual characteristics, and the results are virtually unchanged. The fact that individual characteristics do not have a direct impact on the interest rate, conditional on lender and product, is not surprising given the institutional context described above. The UK mortgage market essentially works as a mortgage supermarket in which a given type of product is offered at a given price, independently of who buys it.

When estimating the interest rate jumps from the coefficients on the LTV bin dummies in equation (1), we are holding all non-LTV mortgage characteristics constant on each side of the LTV threshold. Referring to Figure **??** discussed earlier, if a household is observed in a "5 year fixed standard" mortgage (in a particular bank and month) just below the notch, we are asking how much higher the interest rate would have been for that same product just above the notch. In practice, if the household did move above the notch it might decide to re-optimize in some of

the non-LTV dimensions—say move from 5-year fixed to 2-year fixed rate—and this would give a different interest rate change. However, not only are such interest rate changes endogenous, they are conceptually misleading due to the fact that the non-interest characteristics of the mortgage have value to the borrower and are priced into the offered interest rate. Our approach of conditioning on non-LTV characteristics when estimating the interest rate schedule is based on a no-arbitrage assumption: within a given LTV bin, if lower-interest rate products or banks are available, in equilibrium this must be offset by less favorable terms in other dimensions. In this case, the within-product interest rate jump around the threshold is the right measure of the true price incentive.

3 Bunching Evidence and Empirical Approach

3.1 Actual and Counterfactual Leverage Distributions

Given the interest rate schedule shown in Figure 1, we expect to observe bunching just below notches and missing mass (holes) above notches in the LTV distribution. The idea is to use these empirical moments to identify parameters of interest. Therefore, as a starting point, Figure 2 plots the observed LTV distribution for UK homeowners between 2008-14 around the different notches (depicted with vertical dashed lines). The figure distinguishes between the LTV distribution in the full population of mortgagors (Panel A) and in the sample of remortgagors (Panel B). It is clear that there is very large and sharp bunching below every notch along with missing mass above every notch, consistent with the strong incentives created by the interest rate jumps at those points. The remortgagor distribution is naturally shifted to the left compared to the full distribution due to amortization. For example, there are relatively few households above an LTV of 85% and almost none above 90% in the remortgagor distribution, while there is significant mass at those LTV values in the full distribution. As described above, this paper will focus on the remortgagor sample for which house values are pre-determined, because this allows us to isolate mortgage demand from housing demand.

To precisely estimate the amount of bunching and missing mass in the observed LTV distribution, we need an estimate of the counterfactual LTV distribution—what the distribution would have looked like without interest rate notches—and the public finance literature has developed approaches to obtain such counterfactuals. The standard approach is to fit a flexible polynomial to the observed distribution, excluding data around the notch, and then extrapolate the fitted dis-

tribution to the notch (Chetty *et al.* 2011; Kleven & Waseem 2013). However, this approach is not well-suited for our context: it it based on the assumption that notches affect the distribution only locally, which may be a reasonable assumption when there is only one notch or if the different notches are located very far apart. This is not satisfied in our setting in which we have many notches located relatively close to each other, and where Figure 2 suggests that all parts of the distribution are affected by notches. For example, it would be difficult to evaluate the counterfactual density at the 75% LTV notch using observations further down the distribution, say around 70%, because those observations are distorted by other notches.

To resolve this issue, we propose a new approach to assess the counterfactual distribution that exploits the panel structure of the remortgagor data. Based on the LTV in the previous mortgage, the amortization schedule, and the house value at the time of remortgage (which is assessed by the lender and therefore outside the remortgagor's control), we know the new LTV *before the remortgagor has taken any action*. We label this the *passive* LTV which would be the LTV in the new mortgage if the homeowner simply rolled over the outstanding debt from the existing mortgage. We will base our estimate of the counterfactual LTV on the passive LTV with an adjustment that we describe below.

In Panel A of Figure 3 we compare the actual LTV distribution for remortgagors—i.e., the observed distribution shown in the previous figure—to the passive LTV distribution. We see that the passive LTV distribution is smooth overall; unlike the actual LTV distribution it features no excess bunching or missing mass around notches. In general, the two distributions in Figure 3A may be different for two reasons: (i) behavioral responses to notches, and (ii) equity extraction or injection that would have happened even without notches. The second effect does not create bunching or missing mass, but it may smoothly shift the distribution. In this case, the passive LTV distribution would not exactly capture the counterfactual LTV distribution. To gauge the importance of such effects, we use information on equity extracted among households who do *not* bunch at notches in the final LTV distribution. Figure A.3 shows that equity extracted among non-bunchers is positive through most of the passive LTV distribution (except at the very top) and has a smooth declining profile. We adjust the passive LTV distribution for non-bunching effects on leverage using the profile of equity extracted in Figure A.3. The assumption we are making is that this equity extraction profile (for non-bunchers) is a good proxy for the entire population of remortgagors (including bunchers) in the counterfactual scenario without notches. Given that equity extracted is in general positive, this adjustment shifts the passive LTV distribution to the right and gives the counterfac-

tual LTV distribution shown in Figure 3B. The adjustment has a small effect due to the fact that the passive LTV distribution is very flat, and this suggests that more sophisticated procedures will not have significant effects either.[7]

Comparing the actual and counterfactual LTV distributions in Panel B provides clear visual evidence of bunching and missing mass around each notch. Notice that, except for the region below the bottom notch at 60%, the actual and counterfactual distributions never line up. This is because the actual distribution below each notch is affected by missing mass due to a notch further down. This implies that the standard approach to obtaining the counterfactual—fitting a polynomial to the observed distribution, excluding data right around the notch—would produce biased estimates in our context.

Our measure of the counterfactual distribution has three key advantages. First, it is a *global* counterfactual distribution as opposed to the local counterfactuals in previous work. That is, we are able to say how interest rate notches affect the entire leverage distribution. Second, a specific challenge in the previous literature has been how to obtain a counterfactual in the presence of extensive margin responses, i.e. the fact that the notch may make agents drop out of the market entirely (Kleven & Waseem 2013; Kopczuk & Munroe 2014). Extensive responses do not seem very realistic here—i.e., existing homeowners who instead of re-mortgaging drop out of the market entirely due to the interest rate notches—but even if they were, we are not sensitive to such responses. This is because we have a counterfactual LTV for each individual, conditional on remortgaging, and can therefore cleanly measure the intensive margin response. Third, the counterfactual distribution takes as given that the re-mortgagors also faced notches in their previous mortgages and were potentially bunching in response to those. That is, the counterfactual distribution is generated from previous LTVs (that featured bunching) combined with idiosyncratic amortization and house price changes. This is not a problem for our strategy, but is in fact exactly the right counterfactual experiment: we want to estimate how people are responding to a notched incentive at time t taking as given past behavioral responses to notches they faced prior to time t. For all of these reasons we view the counterfactual shown in Figure 3 as very compelling.

[7] A more sophisticated approach to adjusting for counterfactual equity extraction among bunchers would be an iterative procedure that identifies a fixpoint: the assumed counterfactual equity extraction for bunchers at notch n, X_n, should be such that the estimated amount of bunching implies counterfactual equity extraction for bunchers at n equal to X_n. Again, this is unlikely to make much of a difference in our setting, because we are shifting a distribution that is very flat.

3.2 Conceptual Approach

By comparing the observed and counterfactual LTV distributions shown above, we can obtain estimates of bunching at notches that allow us to evaluate LTV responses to interest rate changes. Such reduced-form estimates are valuable in their own right as they provide some of the first non-parametrically identified evidence of mortgage demand responses to the cost of borrowing. Moreover, by specifying a dynamic model we show that the reduced-form estimates can be linked to a key structural parameter, the Elasticity of Intertemporal Substitution (EIS). Our paper therefore has the dual goal of providing reduced-form estimates of mortgage demand elasticities (without parametric assumptions) as well as structural estimates of the EIS (making parametric assumptions). While we describe the intertemporal theory that underlies this estimation in detail in section 5, here we provide a more intuitive outline of the conceptual approach.

We denote the LTV by λ and consumption today and in the future, respectively, by c_0 and c_1. Our approach consists of the following steps:

- Given the intertemporal budget constraint, the interest rate notch at an LTV threshold λ^* translates into a notch at a consumption threshold c_0^*.
- Given preferences, the notch in consumption leads to intertemporal substitution responses $\Delta c_0, \Delta c_1$ that relate to the EIS. Using the budget constraint, this relates the leverage response $\Delta\lambda$ to the EIS.
- The leverage response maps into an amount of bunching in the LTV distribution, and hence this links bunching to the EIS.

The approach is illustrated in the intertemporal budget set diagram in Figure 4. In the absence of an interest rate notch, the intertemporal budget set is given by the linear curve, the slope of which equals minus the gross interest rate, $-R$. The introduction of an interest rate notch at an LTV threshold λ^* creates a notch at a corresponding consumption threshold c_0^*: as current consumption exceeds this threshold, the interest rate jumps discretely to $R + \Delta R$ and so future consumption drops discretely. This notch induces homeowners (remortgagors) who are initially consuming just above the cutoff c_0^* to reduce their current consumption to the threshold, thereby getting larger consumption in the future. The degree to which remortgagors are willing to engage in such intertemporal substitution depends on the curvature of indifference curves, which is determined by the value of the EIS. The *marginal bunching household* is illustrated in the diagram: this household is

exactly indifferent between the notch point c_0^* and its best interior location above the notch c_0^I. The reduction in current consumption for this household equals Δc_0 and, using the budget constraint and information on the value of the house, this corresponds to an LTV response of $\Delta\lambda$. All households initially located between the LTV threshold λ^* and the marginal buncher at $\lambda^* + \Delta\lambda$ move down to the notch point. Denoting the counterfactual LTV distribution by $f_0(\lambda)$, this produces bunching equal to $B = \int_{\lambda^*}^{\lambda^*+\Delta\lambda} f_0(\lambda)\, d\lambda \approx f_0(\lambda^*)\,\Delta\lambda$.

The fundamental idea of our approach—a dynamic extension of the static framework by Kleven & Waseem (2013)—is that we can use the indifference condition between the notch point c_0^* and the best interior position c_0^I for the marginal buncher to derive a condition that relates the EIS to empirical entities. The approach works as follows. Denoting the EIS by σ and a vector of household characteristics by X (including income, passive LTV, LTV at the notch, house value, current and future interest rates, and the discount factor), it is possible to write the lifetime utility of the marginal buncher at the notch as $U^N(\sigma, X)$. Here we exploit Euler equations governing the optimal future consumption path conditional on bunching at the notch today. Furthermore, the lifetime utility of the marginal buncher at the interior can be written as $U^I(\sigma, \Delta\ln\lambda, \Delta\ln R, X)$, where we use Euler equations that have to hold in the interior. The marginal buncher is indifferent between these two options and therefore

$$F(\sigma, \Delta\ln\lambda, \Delta\ln R, X) \equiv U^N(\sigma, X) - U^I(\sigma, \Delta\ln\lambda, \Delta\ln R, X) = 0 \tag{2}$$

$F(\cdot) = 0$ is an equation that can be solved for the single unknown parameter σ given the estimated value of $\Delta\ln\lambda$ (which we obtain from bunching as discussed above) and the observed/calibrated values of $\Delta\ln R$ and X. The vector X includes two parameters that we do not directly observe and therefore have to calibrate (the discount factor and the path of future interest rates), but we will show that these have virtually no effect on the estimation. The reason is that the X parameters scale utility at the notch and at the interior by roughly the same amounts and approximately cancel out in the estimating indifference equation (2).

The exposition above implicitly assumes that there is just one value of the structural elasticity σ, while in practice there is likely to be heterogeneity in this parameter. In fact, the empirical LTV distributions shown in section 3.1 suggests that this has to be the case: without heterogeneity, there would be an empty hole in the LTV distribution between λ^* and $\lambda^* + \Delta\lambda$ (the width of which depends on the size of σ), but the empirical distributions feature a more diffuse missing mass and

some remortgagors are located just above the notch.[8] This provides *prima facie* evidence that some households have very small σs while others have larger σs. In this case bunching identifies the average response $E[\Delta\lambda]$ and the structural elasticity in equation (2) will be evaluated at the average response (see Kleven & Waseem 2013).[9] We will directly explore heterogeneity by estimating the average response $E[\Delta\lambda]$ in different subsamples using our rich data.

4 Mortgage Demand Responses to Interest Rates: Reduced-Form Estimation

In this section we present non-parametric evidence on the amount of bunching at LTV notches and the implied elasticities of mortgage demand with respect to the interest rate. Despite the importance of such mortgage demand elasticities to policy questions in household finance, public finance and macro, there is virtually no evidence on what might be a reasonable magnitude for them. These mortgage demand elasticities are not structural parameters, but the analysis here will be a stepping stone towards our structural estimation.

We first consider all notches together by pooling our data into a single average notch, which allows us to analyze average behavior in the population. Specifically, for each notch point n and each mortgage i, we calculate a normalized LTV as $LTV_{in} = LTV_i - n$.[10] We then stack the normalized LTVs across the five notches and consider their distribution around the average notch at zero. This is shown in Figure 5 in which the actual LTV distribution is the series in black dots and the counterfactual LTV distribution is the series in orange crosses, with their frequencies in 0.25%-point bins given on the left axis. The counterfactual is obtained using the panel method described in Section 3.1, with the global distribution shown in Figure 3. The figure also plots the pooled conditional interest rate at each LTV in green squares (right axis) obtained from the non-parametric regression described in Section 2.3.[11] As shown in the figure, the interest rate jump at the average notch equals

[8] Besides very small σs among some households, the presence of density mass just above the notch may reflect various optimization frictions (including liquidity constraints), an issue that we will directly address when estimating the structural elasticity as described in section 5.

[9] To be precise, $E[\Delta\lambda]$ captures the average response across *marginal bunchers* with different σs. This is in contrast to the average response by *all bunchers*, including inframarginal bunchers who are strictly better off at the notch and whose response therefore does not pin down σ.

[10] For example, a mortgage with an LTV of 73% has normalized LTVs of $LTV_{i70} = 3$ with respect to the 70% notch and $LTV_{i75} = -2$ with respect to the 75% notch.

[11] The pooled interest rate line is a weighted average of conditional interest rates from the regression (1) performed separately around each notch. The weights for each notch are given by the number of mortgages in the bin containing the notch in the counterfactual distribution.

$\Delta r = 0.25\%$-points relative to a baseline of $r = 3.35\%$ below the average notch. The figure also shows estimates of bunching scaled by the counterfactual density at the threshold, $b \equiv B / f_0(\lambda^*)$, and the size of the LTV response $\Delta\lambda$ obtained from the equation $B = \int_{\lambda^*}^{\lambda^*+\Delta\lambda} f_0(\lambda)\, d\lambda$ as explained in the previous section. We measure bunching B as excess mass in the actual distribution relative to the counterfactual distribution in a 1%-point range below the threshold, allowing for some overshooting among bunching households. However, because bunching is extremely sharp in our setting, the results are very robust to reducing the width of this range. Finally, standard errors are computed by bootstrapping, drawing 100 samples of mortgages from the observed sample (with replacement, stratifying by notch) and re-estimating interest rate jumps, bunching, LTV responses, etc. in each bootstrap iteration.

Several key findings emerge from Figure 5. First, there is very large and sharp bunching equal to almost six times the height of the counterfactual distribution at the notch. Second, there is a very clear gap between the actual and counterfactual distributions to the right of the notch; this is the "hole" where the bunching households would have been observed absent the notched interest rate schedule. The hole extends to 3-4%-points above the notch, implying that the *most* responsive households reduce their LTV by almost 4%-points in response to the interest rate notches. Third, the average LTV response to interest rate notches is equal to 1.74%-points. As defined above, the average response corresponds to how far one has to go in the counterfactual distribution above the notch to account for the excess mass at the notch. Note that this is a smaller number than the width of the hole (about 4%-points), which reflects the entire range in which some households are responding. Fourth, unlike previous bunching approaches that estimate counterfactuals by fitting polynomials, our observed counterfactual distribution does not line up with the actual distribution below the bunching area. This is because our counterfactual accounts for the global LTV distortions created by the presence of multiple notches, which the polynomial approach cannot address.

Fifth, the graph shows that there is a non-trivial number of mortgages observed just above the notch point. Missing mass in the first bin above the notch is about 70%, significantly lower than the 100% one might have expected given the enormous incentive to reduce LTV for households located so close to the notch. The standard interpretation of incomplete holes above notches—especially in the context of labor supply and taxable income responses—has been the presence of *optimization frictions* such as switching costs, inattention, and misperception (Kleven & Waseem 2013). For example, remortgagors may face real or perceived costs of switching to a different bank or contract that provides a better deal below the notch. It is also possible that the complex intertemporal in-

centives associated with choosing a mortgage are not fully understood by all homeowners. Some households may not fully understand the amortization schedule of the loan—how monthly payments are split between principal and interest over the life of the loan—and instead focus simply on the total monthly payment. Such misperception will attenuate responses to interest rate jumps if some of the extra interest payments are absorbed by smaller principal payments, leading to a smaller increase in the total monthly payment. Finally, while it is conceivable that some borrowers stay above notches due to liquidity constraints that prevent them from moving down the LTV distribution, the patterns of equity extracted in our data suggest that this effect is not first order here. We saw in Figure A.3 that equity extracted among non-bunchers is positive and sizable through most of the LTV distribution, implying that it would be feasible for most of them—especially for those located just above an LTV notch—to move below the notch by taking out less equity. Specifically, households located in the first bin above a notch would need less than 0.25% of their house value in order to bunch, which is smaller than the equity extracted for the vast majority of our sample. In the next section, we return to the question of optimization frictions and how we deal with it in the structural estimations.

Figure 6 shows bunching evidence for each individual notch, but is otherwise constructed in the same way as the previous figure (we do not show the 60% notch, but it is qualitatively similar to the ones displayed). The interest rate jumps shown in this figure are somewhat smaller than those reported in Figure 1, because the interest rate jumps here apply to the sample of remortgagors as opposed to the full sample including first-time mortgagors.[12] The evidence from the individual notches is qualitatively consistent with the evidence from the pooled notch; excess bunching is large and leaves a hole in the LTV distribution above each notch. The degree of bunching b and the leverage response ΔLTV are larger at higher LTVs, but so are the interest rate jumps and thus the incentives to respond.

Based on the bunching evidence, we now turn to the estimation of reduced-form mortgage demand elasticities in Table 2. The table shows results for the five individuals notches (60, 70, 75, 80, 85) and for the pooled average notch. Panel A of the table summarizes the statistics presented so far: the interest rate below each notch r, the interest rate jump at each notch Δr , excess bunching b, and the leverage response ΔLTV. To provide context, Panel B of the table presents the mean value of four descriptive statistics in the bunching range below each notch: loan size L^*, house value V^*,

[12] UK banks tend to impose larger interest rate jumps on first-time mortgagors at high LTV ratios, which explains why the difference is most pronounced at the 80% and 85% LTV notches.

time between mortgages Δt, and equity extracted in proportion to the house value $\Delta Equity/V$. On average, bunchers have about £170,000 of mortgage debt, refinance every third year, and extract substantial amounts of home equity when refinancing. As we discuss later, equity extracted is strongly related to house price appreciation since the last mortgage, a pattern that is consistent with findings for the US by Mian & Sufi (2011).

Panel C presents four key statistics: the amount of current consumption foregone due to LTV responses (Δc_0), the amount of future consumption gained due to LTV responses (Δc_1), the implied rate of return on reducing the LTV (r^* as defined below), and the reduced-form elasticity of mortgage demand with respect to the gross marginal interest rate (ε). The reduction in current consumption is calculated as $\Delta c_0 = \Delta\lambda \times V$ and is equal to about £4,000 on average, ranging from £1,652 at the bottom notch to £8,210 at the top notch. The increase in future consumption Δc_1 is created by two effects: lower interest payments on inframarginal units of debt when moving below the notch ($\Delta r \times \lambda^* \times V$) and lower interest payments on the marginal units of debt ($(r + \Delta r) \times \Delta\lambda \times V$) until the next re-mortgage event. The future consumption gains are very large: on average these consumption gains equal about £6,000, 50% higher than the current consumption loss, and they range from £2,524 at the bottom notch to £11,895 at the top notch.

Using our estimates of Δc_0 and Δc_1, it is possible to calculate the (annualized) marginal rate of return on reducing the LTV around the notch as $r^* \equiv \left(\frac{\Delta c_1}{\Delta c_0}\right)^{1/\Delta t} - 1$. Given that the underlying LTV and consumption responses are based on the concept of the *marginal* bunching household (as explained in section 3.2), this should be interpreted as the rate of return obtained by a marginal buncher who reduces leverage by $\Delta\lambda$ as shown in the table.[13] Inframarginal bunchers—those located between the threshold λ^* and the marginal buncher at $\lambda^* + \Delta\lambda$—obtain larger rates of return from bunching, with the rate of return converging to infinity when approaching the notch from above. The table shows that the marginal rates of return r^* are large, about 11% on average and exceeding 7% at any LTV notch. When viewing these estimates as the return foregone by the marginal households who are *not* bunching, their magnitude raise a financial puzzle comparable to the widely discussed equity premium puzzle (Mehra & Prescott 1985; Kocherlakota 1996; Siegel & Thaler 1997). Importantly, because paying down mortgage debt to get below the interest rate notch is essentially a risk-free investment opportunity, risk aversion cannot be central to explaining this puzzle. More natural candidates are the optimization frictions discussed above.

[13] More precisely, in a world with heterogeneity in underlying responsiveness, bunching estimates represent the *average* response across *marginal* bunchers that vary with respect to their underlying elasticities (Saez 2010; Kleven & Waseem 2013).

As the final piece of the reduced-form analysis, we present estimates of mortgage demand elasticities. Based on the approach developed by Kleven & Waseem (2013) in the context of tax and labor supply, we define the reduced-form elasticity of mortgage demand with respect to the gross interest rate as follows

$$\varepsilon \equiv \frac{\Delta\lambda/\lambda^*}{\left(R^* - R\right)/R} \tag{3}$$

where $R^* \equiv 1 + r^*$ is the gross marginal rate of return and $R \equiv 1 + r$ is the gross interest rate below the notch. As shown by Kleven & Waseem (2013), such reduced-form elasticities obtained from notches are approximations; the discontinuous nature of notches necessitates a model to estimate exact elasticities (as we do in the next section).

The reduced-form mortgage elasticities are sizable and feature substantial heterogeneity with respect to leverage. The average elasticity obtained from the pooled notch equals 0.25. The elasticity is increasing in leverage, ranging from just below 0.1 at the 60% LTV notch to about 1.4 at the 85% LTV notch. These elasticities suggest that changes in the net-of-tax interest rate—for example through monetary policy or through tax policy—lead to significant changes in household leverage. While these reduced-form estimates are economically important—and represent some of the first evidence of its kind—their interpretation and external validity is potentially limited by the fact that they do not represent any structural parameter. Hence we now turn to the estimation of structural parameters.

5 Elasticity of Intertemporal Substitution: Structural Estimation

Having estimated reduced-form mortgage demand elasticities, we now specify a theoretical model that allows us to estimate a key structural parameter for a range of questions in economics: the elasticity of intertemporal substitution. This requires stronger assumptions, but has the benefit of providing a deeper and more externally valid parameter. What is more, we show that our estimates are robust to a large range of parametric assumptions.

5.1 Theoretical Framework

SETUP: We model the consumption and debt choices of remortgagors, assuming an infinite horizon and perfect foresight. Each household, indexed by i, has preferences $U\left(C^i, H^i\right)$ over lifetime

consumption C^i and real housing wealth H^i. Consumption C^i takes the CES/CRRA form

$$C^i \equiv \left(\sum_{t=0}^{\infty} \delta^t \left(c_t^i \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{4}$$

where σ is the elasticity of intertemporal substitution, our main parameter of interest, and δ is the discount factor. Focusing on remortgagors, we assume that housing wealth H^i was chosen prior to period 0 and represents the steady state level of housing that homeowners desire in all future periods.[14] Households may refinance their house in any period t, but we focus specifically on the remortgage decision in period 0 allowing for an optimal future path. The key assumption on preferences is that housing consumption is weakly separable from lifetime (non-housing) consumption, which simplifies the analysis considerably.

Household i enters period 0 with a mortgage debt of $R_0^i D_0^i$. The household chooses mortgage debt D_{t+1}^i in each period t, incurring a gross interest rate of R_{t+1}^i on this debt. Given housing prices P_t, the choice of debt implies an LTV of $\lambda_{t+1}^i \equiv \frac{D_{t+1}^i}{P_t H^i}$.

The household obtains a stream of income y_t^i, so that its budget constraint in period t reads

$$c_t^i \leq y_t^i - R_t^i D_t^i + D_{t+1}^i. \tag{5}$$

This implies a lifetime budget constraint of

$$\sum_{t=0}^{\infty} \frac{c_t^i}{\prod_{s=1}^{t} R_s^i} = \sum_{t=0}^{\infty} \frac{y_t^i}{\prod_{s=1}^{t} R_s^i} - \bar{\lambda}_0^i P_0 H^i, \tag{6}$$

where $\bar{\lambda}_0^i \equiv \frac{R_0^i D_0^i}{P_0 H^i}$ denotes the LTV at the beginning of period 0 prior to any refinancing decision. Note that $\bar{\lambda}_0^i$ corresponds to what we labeled the *passive* LTV in the reduced-form empirical analysis, with the empirical distribution shown in Figure 3A. For a given income stream, the passive LTV determines lifetime consumption possibilities. Consistent with the data, we allow for heterogeneity in $\bar{\lambda}_0^i$ (as well as in housing wealth H^i and income y_t^i).

The budget constraint (6) assumes that households hold no assets other than housing and no liabilities other than mortgage debt. In this case, leverage decisions when remortgaging map into consumption decisions, which is central to our estimation of σ. If households hold non-housing assets, a potential concern is that changes in leverage represent changes in portfolio composition

[14] This steady state assumption can be relaxed, but doing so would require us to take a stand on future growth in housing consumption and prices when when estimating σ.

rather than in consumption. While we do not observe non-housing wealth in the data, we do observe several other household variables that are correlated with wealth such as income and age. It is unlikely that young homeowners with relatively low income have much financial wealth that they can use when refinancing their house (Cloyne & Surico 2013), and so our analysis of heterogeneity in section 5.2 can shed light on the potential role for portfolio composition effects.[15] Importantly, if portfolio composition effects do confound our elasticity estimates to some degree, these estimates will represent upper bounds on the true value of σ. These upper bounds will be very informative as our elasticity estimates are small.

In what follows we analyze how borrowing and consumption decisions in period 0 respond to interest rate notches (and how these responses are linked to the structural parameter σ). In order to characterize the period 0 remortgaging problem, we start by considering the the future path of consumption and debt from period 1 onwards.

Optimization in periods $t \geq 1$: We assume that households will be at interior LTVs in future periods, i.e. households anticipate not to be bunching at LTV notches in the future. Given the weak separability between housing and consumption such interior positions satisfy a standard Euler equation:

$$c_{t+1}^i = \left(\delta R_{t+1}^i\right)^\sigma c_t^i. \tag{7}$$

This allows us to rewrite the consumption aggregate C^i as a function of consumption in periods 0 and 1 alone:

$$C^i = \left(\left(c_0^i\right)^{\frac{\sigma-1}{\sigma}} + \delta \tilde{R} \left(c_1^i\right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{8}$$

where

$$\tilde{R} \equiv \sum_{t=1}^{\infty} \left(\delta^\sigma\right)^{t-1} \prod_{s=1}^{t-1} \left(R_{s+1}^i\right)^{\sigma-1}. \tag{9}$$

A household's borrowing choice in period 0 is a choice between consumption today and consumption in the future. When taking into account that the household will make optimal choices in all future periods, this problem can be represented as a choice between consumption in periods 0 and 1 alone. Equation (8) shows that $\tilde{R}$ is a sufficient statistic for the entire future path of interest rates that determines the relative price between consumption in periods 0 and 1. Any two paths of interest rates for $t > 1$ that lead to the same value of $\tilde{R}$ will lead to the same choice of consumption

[15] Cloyne & Surico (2013) show that the median mortgagor in the UK has zero net financial wealth, limiting the potential importance of portfolio composition effects in our setting.

and borrowing in period 0.[16]

For later use, we note that the lifetime budget constraint from period 1 onwards implies that consumption in period 1 must satisfy

$$c_1^i = \frac{Y^i - R_1^i D_1^i}{\tilde{R}}, \tag{10}$$

where we have used the Euler equation (7) and where $Y^i \equiv \sum_{t=1}^{\infty} y_t^i \prod_{s=1}^{t-1} \left(R_{s+1}^i\right)^{-1}$ denotes the net present value of the household's income from period 1 onwards.

Optimization in period $t = 0$ without notches (counterfactual): For the remortgage decision in period 0, we first consider a counterfactual scenario without notches. In this scenario, each household i can borrow at any leverage ratio at the gross interest rate of $R_1^i = R$. This implies an Euler equation in period 0 given by

$$c_1^i = (\delta R)^{\sigma} c_0^i. \tag{11}$$

This condition harks back to the no-notch tangency point in Figure 4 and gives counterfactual consumption choices in periods 0 and 1. Using equation (10) and the period-0 budget constraint, these consumption choices can be mapped back to a counterfactual choice of borrowing

$$D_1^{iCF} = \frac{\frac{Y^i}{\tilde{R}} + (\delta R)^{\sigma} \left(R_0^i D_0^i - y_0^i\right)}{\frac{R}{\tilde{R}} + (\delta R)^{\sigma}}. \tag{12}$$

This implies a counterfactual LTV ratio of $\lambda_1^{iCF} = \frac{D_1^{iCF}}{P_0 H^i}$, which has the empirical distribution shown in Figure 3B.

Optimization in period $t = 0$ with a notch: We now introduce an interest rate notch at the LTV ratio of λ^*. Households can borrow at $R_1^i = R$ when choosing $\lambda \leq \lambda^*$ and at $R_1^i = R + \Delta R$ when choosing $\lambda > \lambda^*$.

As described in section 3.2, our empirical analysis identifies the LTV response by the marginal buncher who is indifferent between bunching at the notch λ^* and choosing the best interior posi-

[16] As mentioned, the simple representation in (8)-(9) assumes an interior choice of consumption, i.e. that the household does not anticipate to bunch at LTV notches in the future. This assumption skirts the complexity of a strategic choice of debt today in order to increase the likelihood that (after housing housing price changes and amortization) the household will be able and willing to bunch at an LTV threshold in the future. Allowing for such sophisticated dynamic bunching behavior would significantly complicate the model, although it would remain qualitatively similar. Ultimately, the household faces a trade-off between consumption in period 0 and in all future periods and strategic repayment to reach future mortgage notches simply changes the relative weights on current vs. future consumption.

tion above the notch λ^I. The marginal buncher is illustrated in the budget set diagram in Figure 4, and the LTV response by this household is given by $\Delta\lambda = \lambda^{CF} - \lambda^*$. In what follows we use the indifference condition between the notch point and the best interior position to derive a relationship between the estimated LTV response $\Delta\lambda$ and the intertemporal elasticity of substitution σ. Since we now focus on a specific type of remortgagor—the marginal buncher—we suppress the household index i in the expressions below.

We first derive the utility level at the unconstrained, interior position I. Since this is a standard tangency point, the choice of debt at this point D_1^I can be characterized in the same way as counterfactual debt in equation (12), but evaluated at the higher interest rate $R + \Delta R$ above the notch. Based on this debt level and the period-0 budget constraint $c_0^I = y_0 - R_0 D_0 + D_1^I$, consumption in period 0 can be written as

$$c_0^I = \frac{\frac{Y}{\tilde{R}} + \frac{R+\Delta R}{\tilde{R}}\left(y_0 - R_0 D_0\right)}{\frac{R+\Delta R}{\tilde{R}} + \delta^\sigma \left(R + \Delta R\right)^\sigma}. \tag{13}$$

This expression depends on the net present value of future income Y, which is not observed in the data. However, we can eliminate Y by exploiting the fact that household borrowing in the counterfactual scenario D_1^{CF} also depends on future income as can be seen in equation (12). Isolating $Y / \tilde{R}$ in equation (12) and inserting into equation (13), we obtain

$$c_0^I = \frac{D_1^{CF}\left(\frac{R}{\tilde{R}} + (\delta R)^\sigma\right) + \left(\frac{R+\Delta R}{\tilde{R}} + (\delta R)^\sigma\right)\left(y_0 - R_0 D_0\right)}{\frac{R+\Delta R}{\tilde{R}} + \delta^\sigma \left(R + \Delta R\right)^\sigma}. \tag{14}$$

The utility level implied by consumption c_0^I can be derived from the consumption aggregate (8) and the Euler equation (11), which can be rearranged as

$$C^I = c_0^I \left(1 + \delta^\sigma \tilde{R}\left(R + \Delta R\right)^{\sigma-1}\right)^{\frac{\sigma}{\sigma-1}}. \tag{15}$$

Equations (14) & (15) together give consumption utility as a function of entities that we estimate, observe, or can calibrate. By noting that $D_1^{CF} \simeq D_1^* \left(1 + \Delta \ln \lambda\right)$ and $R + \Delta R \simeq \left(1 + \Delta \ln R\right) R$, we can write

$$C^I = C^I\left(\sigma, \Delta \ln \lambda, \Delta \ln R, X\right), \tag{16}$$

where X is a vector that includes $\left(y_0 - R_0 D_0, D_1^*, R, \tilde{R}, \delta\right)$. Apart from σ, the only variables in this expression that are not directly estimated or observed are the discount factor δ and the future path of interest rates, as captured by $\tilde{R}$. We choose reasonable values for these two parameters in

our baseline estimates, but show in Section 5.2 that our results are virtually unchanged for a wide range of values.

We now turn to utility at the notch. From the budget constraint (5), consumption in period 0 when bunching at the notch is given by

$$c_0^* = y_0 - R_0 D_0 + D_1^* \tag{17}$$

From equations (10) and (12), consumption in period 1 when bunching can be written as

$$c_1^* = D_1^{CF}\left(\frac{R}{\tilde{R}} + (\delta R)^\sigma\right) + (\delta R)^\sigma (y_0 - R_0 D_0) - \frac{R}{\tilde{R}} D_1^*. \tag{18}$$

Here we have again used the counterfactual debt choice D_1^{CF} to eliminate future income expectations Y from the expression. Inserting the consumption levels c_0^* and c_1^* into the consumption aggregate (8), we can express consumption utility at the notch as follows

$$C^N = C^N\left(\sigma, \Delta \ln \lambda, X\right), \tag{19}$$

where X includes the same variables as before.

Using the fact that the marginal buncher is indifferent between the notch and the best interior position above the notch, we can characterize the estimating equation as

$$F\left(\sigma, \Delta \ln \lambda, \Delta \ln R, X\right) \equiv C^N\left(\sigma, \Delta \ln \lambda, X\right) - C^I\left(\sigma, \Delta \ln \lambda, \Delta \ln R, X\right) = 0, \tag{20}$$

corresponding to equation (2) shown earlier. This equation allows us to estimate the structural elasticity σ based on our reduced-form bunching estimate $\Delta \ln \lambda$, the interest rate jump at the notch $\Delta \ln R$, and a set of observed and calibrated parameters in X. As mentioned above, the only unobserved variables in X are the discount factor δ and the future interest rate path as reflected in the sufficient statistic $\tilde{R}$. Our benchmark assumption will be that $R_t = R$ in all future periods $t > 1$ in which case we have $\tilde{R} = \frac{1}{1-\delta^\sigma R^{\sigma-1}}$.

We conduct robustness checks for a large range of values of both $\tilde{R}$ and δ in the following section. Estimates of σ are not sensitive to these values. It is worth dwelling on the intuition for the robustness of our results as it highlights an advantage of using notches to estimate this elasticity. Both time preference and interest rate expectations affect borrowing choices by determining the

relative value of future consumption. Changes in these parameters, however, affect the value of future consumption at the notch and at the interior—and therefore the terms C^N and C^I—by similar magnitudes. Their effect on the elasticity σ estimated from equation (20) is thus second order. Put differently, changing the effective rate of time preference through $\tilde{R}$ and δ would affect the level of borrowing and shift the entire LTV distribution, but the response of borrowing to an interest rate notch—a local change in the relative price between consuming today and in the future—is governed primarily by the EIS.

OPTIMIZATION FRICTIONS: The framework presented above assumes that there are no optimization frictions (such as switching costs or misperception). However, as discussed in section 4, households may well be prevented from bunching due to such frictions, in which case our estimate of σ would be downward biased. To deal with this general problem in empirical research, Kleven & Waseem (2013) developed a non-parametric frictions adjustment based on the presence of strictly dominated regions of behavior above notches. In their setting, strictly dominated regions of the income distribution just above tax notches could be used to estimate the fraction of agents that are non-optimizers while being agnostic about the specific reasons for not optimizing. Assuming that the fraction of non-optimizers is the same outside the dominated region (i.e., where it cannot be directly measured), Kleven & Waseem (2013) showed that it is possible to adjust the bunching estimates for the amount of optimization friction in order to estimate true structural elasticities.

Here we propose a parametric version of the Kleven-Waseem friction approach. In our setting, there are no strictly dominated regions per se. Locating just above an LTV notch implies a large drop in future consumption, but allows for larger current consumption. If a consumer is perfectly impatient ($\delta = 0$), locating in such regions will be optimal. However, as long as consumers value future consumption so that $\delta > 0$, there exists no non-negative elasticity of intertemporal substitution ($\sigma \geq 0$) that can justify locating immediately above an LTV threshold. Even with a zero substitution elasticity, the higher interest rate above the notch creates a wealth effect that should make consumers reduce consumption and leverage today, which is inconsistent with locating extremely close to the notch. Hence, for a given value of $\delta > 0$, we can structurally derive a leverage range above the notch that is inconsistent with any $\sigma \geq 0$. The derivation of this dominated region is provided in appendix B. We estimate the fraction of non-optimizers from this region and adjust our estimates of bunching and thus $\Delta \ln \lambda$ that are used in estimating σ. As long as the fraction of non-optimizers in the dominated region is a good proxy for optimization frictions elsewhere (where the amount of friction is unobserved), our adjusted estimates of σ are not biased by

frictions.

5.2 Elasticity of Intertemporal Substitution

In this section we present estimates of the structural elasticity σ based on equation (20), which links the elasticity to the LTV response $\Delta \ln \lambda$, the interest rate jump $\Delta \ln R$, and a set of observed or assumed parameters X. The estimated LTV responses and interest rate jumps at each notch were reported in Table 2. As explained in section 3.2, in the presence of heterogeneity in σ, the LTV response $\Delta \ln \lambda$ represents the average response across different σ-types, and so our estimation of σ capture the EIS for the average household.

The X parameters used in the estimation are summarized in Table A.1. The table shows the average values of current income y_0, outstanding debt at time of refinancing $R_0 D_0$, and new debt after refinancing D_1^* among households in the bunching region below each notch. The table also shows the assumed values of the two X parameters that we do not observe: the discount factor δ (set equal to 0.95 in the baseline) and expected future interest rates for households at different LTV notches (set equal to the current interest rate they face below the notch). To convert the nominal interest rates in the table into real interest rates, we deduct an inflation rate of 2%, corresponding to the Bank of England's inflation target. Importantly, as we show later, the assumptions we make about the discount factor and future real interest rates have virtually no impact on the estimations. The length of a period Δt is defined as the duration of the initial interest rate (typically 2, 3 or 5 years)—corresponding to the time between remortgaging for most households—and is equal to 3.34 years on average. We convert annual values into this period length.

The estimation of σ is shown in Table 3. As above, standard errors are obtained by block bootstrapping with replacement. The first row of the table restates our estimates of bunching b, as reported in table 2. Using the friction methodology described above, the second row shows the fraction of borrowers a who face some form of optimization friction. For each notch, this is calculated as the number of borrowers observed in the dominated LTV region above the notch as a fraction of the counterfactual number of borrowers in this region. Optimization frictions are substantial: across all notches, more than a third of borrowers do not move down to the notch even though they are located in an LTV region that is dominated by the notch for any $\sigma \geq 0$. The fraction of non-optimizers is particularly large at the 60% notch, which is much smaller than the other notches. This is consistent with a fixed cost model of optimization: households optimize only when the stakes are sufficiently large. Having estimated the fraction of non-optimizers at

each notch, we can re-estimate the amount of bunching that would be observed if only optimizers were present in the data. This is equal to $b_{Adj} = \frac{b}{1-a}$, reported in the third row of the table. Using the adjusted bunching mass, we re-estimate the LTV response $\Delta\lambda$ for optimizers, reported in the fourth row. The average response at the pooled notch is 2.85%-points, compared to 1.74%-points without the friction adjustment as shown in table 2.

Finally, in the last row of Table 3, we report our estimates of the EIS σ based on equation (20). The EIS is small, ranging from 0.05 to 0.23 across the different LTV notches, and is precisely estimated. The average EIS based on the pooled notch is just below 0.1. Note that the structural EIS is in general much lower and features less heterogeneity across LTV levels than the reduced-form debt elasticity shown in Table 2. This result echoes insights from early simulation studies of the life-cycle model showing that, under most calibrations, any given value of the EIS translates into a much larger savings elasticity with respect to the after-tax interest rate (Summers 1981; Evans 1983).[17]

Having observed some heterogeneity in the EIS across LTV ratios, we turn to other dimensions of potential heterogeneity in Table 4. In this table we analyze how our EIS estimate varies with age, income, loan-to-income (LTI), house price appreciation since the previous mortgage, the interest rate change since the previous mortgage (assuming passive borrower behavior), and interactions between age and all the other covariates. For each covariate, we estimate σ in four quartiles of the distribution of the covariate, pooling across the five notches.

The overall finding from Table 4 is that there are no dramatic patterns of heterogeneity. The EIS is relatively stable in most dimensions and our estimates are always below 0.3, a modest value compared to most previous estimates and to benchmark assumptions in calibration studies. Specifically, we find that the EIS is weakly increasing in age, quite stable across income groups, and weakly increasing in the LTI ratio. Interacting age with income and LTI, we find larger elasticities in the top quartile, around 0.2. Given that age, income and LTI are highly correlated with household wealth (which we do not observe in the data), these findings suggest that heterogeneity with respect to wealth is not very strong either. Wealth and income are often argued to be important sources of EIS heterogeneity (see Guvenen 2006), but our evidence does not provide support for this channel being very important.

[17] Summers (1981) showed that, under his calibration, a structural EIS of 0.33 implies a savings elasticity of almost 3, while Evans (1983) showed that the relationship between the structural EIS and the savings elasticity depends critically on other parameters of the environment (although most of his calibrations imply a savings elasticity that is considerably larger than the EIS).

It is interesting to view these heterogeneity results in the context of two potential confounders discussed above: portfolio composition effects (which would inflate our EIS estimates, other things equal) and liquidity constraints (which would attenuate our EIS estimates). We have argued that liquidity constraints are unlikely to have large effects in our setting due to the fact that most households extract enough equity that it would be feasible for them to move down to a notch, but we cannot rule our that such constraints play a role for some households. Both liquidity constraints and portfolio composition effects lead to the prediction that the observed EIS should be larger among households with larger non-housing assets (stock, bonds, cash deposits, etc.), including assets that can be liquidated when remortgaging. Given that age, income and LTI are correlated with financial wealth, the absence of strong heterogeneity in these dimensions limit the potential importance of these potential confounders.

The strongest source of heterogeneity in our data is house price appreciation. Borrowers in the top quartile of the distribution of house price growth have an EIS of 0.24, three times larger than the full sample. Interestingly, house price growth is also strongly related to equity extracted, which increases from 2.4% of the house value in the bottom quartile of house price growth to 5.9% in the top quartile. The strong relationship between house price increases and equity-based borrowing is consistent with the findings for the US by Mian & Sufi (2011). Hence, we find that homeowners who experience large house price increases extract more equity when refinancing, are more likely to bunch at LTV notches, and as a result feature a larger EIS. It is unlikely that this pattern is driven by liquidity constraints, because even in the third quartile of house price growth (where the EIS is still very low) equity extraction is large enough to make it feasible for most borrowers to move down to a notch by extracting less equity. It is possible that house price growth acts as a proxy for deeper EIS determinants. Large house price growth is strongly related to being located in London, which is very different housing market (one that was hardly affected by the recession, see Best & Kleven 2015) and consist of a very different set of borrowers.

5.3 Robustness and Extensions

In estimating the EIS we made a number of modeling assumptions and assigned values to some unobserved parameters. We now describe a number of robustness checks, the results of which are reported in Table 5, showing that our estimates are virtually unchanged under a variety of alternative assumptions. Specifically, our EIS estimates are robust to other assumptions on the path of future interest rates, the discount rate, hyperbolic discounting, risk aversion and uncertainty.

The table shows results for the pooled average notch, but the robustness of our estimates also holds for each notch separately.

FUTURE INTEREST RATES: We have made an admittedly crude assumption on future interest rates, namely that they stay constant at their current level in all future periods. Such beliefs seem unrealistic given the historically low interest rates during the period we study. In Panel A of Table 5 we compare the EIS estimates obtained when varying the future interest rate between 2% and 15% (in nominal terms), with the baseline assumption of unchanged interest rates shown as $r_1 = 3.35\%$. The real interest rates are 2%-points lower as we continue to deduct the Bank of England's inflation target of 2% from nominal interest rates. We keep the assumption of constant future interest rates, but this entails no loss of generality given that future interest rates affect our estimates only through the sufficient statistic $\tilde{R}$ defined in equation (9). Alternative time-varying paths of future interest rates leading to the same value of $\tilde{R}$ would yield the same estimate of σ. In the table we report the implied value of $\tilde{R}$ for each estimation. The table shows that the estimates of σ are very robust to future interest rates, with the elasticity varying from 0.074 to 0.091 over the large range of interest rates we consider, despite the large changes in $\tilde{R}$ across the different scenarios.

DISCOUNT FACTOR AND HYPERBOLIC DISCOUNTING: In Panel B of Table 5 we vary the discount factor δ between 0.7 and 0.99. We find that our results are extremely robust to the discount rate, with the estimated σ varying from 0.081 to 0.076 over the large range of discount factors considered.

Our life-cycle model makes the standard assumption that future periods are discounted at a constant geometric rate δ. Laibson (1997) and others have argued that hyperbolic discounting is more consistent with the experimental evidence on intertemporal choice. We therefore re-estimate our model allowing for hyperbolic discounting using the (β, δ) model. We allow for different degrees of present bias β (keeping the baseline assumption of $\delta = 0.95$), with our baseline model corresponding to $\beta = 1$. The results in Panel C show that even very strong degrees of present bias barely move our estimates. What is the intuition for the irrelevance of present bias in our estimation of σ? While the reluctance of many households in our sample to take advantage of the lower interest rate below the closest notch may seem "myopic" at first glance, the rate of time preference (whether geometric or hyperbolic) plays only a secondary role in this choice. Discount rates reflect mainly in the desired *level* of the LTV, whereas the *response* of the LTV to the local price change created by the notch is governed by the curvature of preferences as captured by the

EIS. This distinction between level and curvature effects is also the reason why both the standard geometric discount factor and future interest rates have very small impacts on our estimations.

Uncertainty and Risk Aversion: Our baseline model assumes perfect foresight, sidestepping the question of uncertainty. Moreover, the CRRA functional form imposes that the coefficient of relative risk aversion is equal to the reciprocal of the EIS. To explore the role of uncertainty and risk aversion, we extend the model to a more flexible functional form that separates risk aversion and the EIS. Specifically, we follow the recursive preference specification of Epstein & Zin (1989) and Weil (1990):

$$C_t = \frac{\left((1-\beta)\, c_t^{\frac{\sigma-1}{\sigma}} + \beta \left[1 + (1-\beta)\,(1-\gamma)\,(E_t\left[C_{t+1}\right]) \right]^{\frac{1-\frac{1}{\sigma}}{1-\gamma}} \right)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}} - 1}{(1-\beta)\,(1-\gamma)}. \tag{21}$$

Here γ is the coefficient of relative risk aversion over static gambles, while σ is the elasticity of intertemporal substitution as before. In Appendix C we provide an analytical solution to the estimating indifference equation under the preference specification (21). To simplify, this characterization assumes that future interest rates R_t are independently and identically distributed in each period, while the path of future income y_t remains certain. Under these assumptions, the indifference equation is shown to be identical to the one derived in Section 5.1, with the modification that $\tilde{R}$ is calculated with respect to the certainty-equivalent expected interest rate $\hat{R}_{t+1} \equiv \left(E_t \left\{ R_{t+2}^{1-\gamma} \right\} \right)^{\frac{1}{1-\gamma}}$. Both uncertainty and the coefficient of risk aversion affect the estimated EIS only through the value of $\tilde{R}$.

The estimation results with Epstein-Zin-Weil preferences are shown in Panel D of Table 5, where we vary the coefficient of risk aversion γ between 0 and 100. In these estimations we allow for large uncertainty over interest rates: the four interest rate levels considered in Panel A (ranging from 2% to 15%) are equally likely in each future period. As we can see from the table, the coefficient of risk aversion has virtually no impact on the EIS.

While these results rely on a number of assumptions, they highlight a more general insight. Choosing whether or not to bunch is a risk-free proposition: Failing to bunch provides extra consumption today, while bunching provides a certain reduction in future interest payments. The degree of risk aversion merely affects the pricing kernel at which these streams of payments are valued. Future income and interest rates, while uncertain, are not affected by the bunching choice and thus provide the same stream of financial flows in both the C^N- and the C^I-terms of the indif-

ference equation. These financial flows are not separable in utility from the choice of LTV, but this non-separability is second-order in the bunching decision.

6 Conclusion

Recent years have seen a surge of interest in household leverage, especially homeowner leverage, and its role for the macro economy and the recent financial crisis. While there has been work on the response of homeowner borrowing to credit supply and house prices (Mian & Sufi 2009, 2011), the response to interest rates has remained an elusive topic for economists due to the difficulties of finding plausibly exogenous variation. This is unfortunate given that the after-tax interest rate is one of the main determinants of household debt that can be directly affected by policy makers through monetary policy, tax policy, regulatory policies, etc. In this paper we have taken a step to fill this gap in knowledge based on a novel source of quasi-experimental variation in interest rates coming from mortgage contract notches in the UK.

We have highlighted that mortgage debt responses to interest rates are related to a more fundamental structural parameter: the Elasticity of Intertemporal Substitution (EIS). The EIS is arguably one of the most important parameters in economics as it plays a central role in almost any economic model involving intertemporal choice. It governs consumption and savings responses to interest rate changes, affects the reaction of consumption to income shocks, is an important parameter for asset pricing, and provides a key statistic for evaluating a range of macroeconomic and microeconomic policies. Despite its importance, there exists no empirical consensus on a reasonable range for this parameter due to limitations in data and research designs. Empirical estimates range between 0 and 2 depending on the setting and empirical design, which is an an enormous range in terms of its policy implications and allows calibration studies too much freedom to be truly informative. We find that the EIS is small, around 0.1-0.2. This is close to the level obtained in the early macro-based literature (Hall 1988), although we have arrived at this conclusion using a fundamentally different approach. Such low values of the EIS have important implications for the extent to which business cycles can be propagated through intertemporal substitution in consumption as posited by all standard business cycle models.

A potential limitation of our research design is that we cannot precisely measure the level of consumption, because we do not observe all the elements of the intertemporal budget constraint. Specifically, while we observe income, housing wealth and mortgage debt, we do not observe non-

housing wealth such as stock holdings and cash deposits. This leaves open the possibility that the intertemporal substitution effects we estimate are driven partly by portfolio composition effects rather than by consumption reallocation. While this would have been an important concern had we found large values of the EIS, it is not a qualitatively important issue given that we find small elasticities. The presence of portfolio composition effects imply that our estimates represent upper bounds on the true EIS and this serves to reinforce our main conclusions. This is consistent with the fact that we find even lower elasticities, around 0.05-0.1, among households that are the least likely to have much non-housing wealth such as younger, lower-income households.

Finally, in this paper we have focused on remortgagors in order to separate mortgage demand responses from housing demand responses. The responsiveness of housing demand to interest rates is an interesting question in its own right, and studying this empirically would be necessary to paint a complete picture of how interest rate changes affect the housing market and the economy. By combining data on remortgagors (who can respond on the debt margin) with data on home buyers (who can respond on both the debt and housing margins), it is possible to develop a strategy that separately identifies the impact of interest rates on debt and housing. The fact that bunching is much stronger among home buyers than among remortgagors (as we can infer from Figure 2) suggests that there *is* a response on the housing margin as well. We leave the detailed study of such responses for future research.

References

AGARWAL, SUMIT, LIU, CHUNLIN, & SOULELES, NICHOLAS S. 2007. The Reaction of Consumer Spending and Debt to Tax Rebates – Evidence From Consumer Credit Data. *Journal of Political Economy*, **115**(6), 986–1019. 4

ATTANASIO, ORAZIO P., & WEBER, GUGLIELMO. 1993. Consumption Growth, the Interest Rate and Aggregation. *Review of Economic Studies*, **60**(3), 631–649. 4

ATTANASIO, ORAZIO P., & WEBER, GUGLIELMO. 1995. Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, **103**(6), 1121–1157. 4

ATTANASIO, ORAZIO P., & WEBER, GUGLIELMO. 2010. Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy. *Journal of Economic Literature*, **48**, 693–751. 4

ATTANASIO, ORAZIO P., GOLDBERG, PINELOPI K., & KYRIAZIDOU, EKATERINI. 2008. Credit Constraints in the Market for Consumer Durables: Evidence From Micro Data on Car Loans. *International Economic Review*, **49**(2), 401–436. 3

AUCLERT, ADRIEN. 2015. *Monetary Policy and the Redistribution Channel*. Mimeo: MIT. 4

BERNHEIM, B. DOUGLAS. 2002. Taxation and Saving. *Chap. 18, pages 1173–1249 of:* AUERBACH, ALAN J., & FELDSTEIN, MARTIN (eds), *Handbook of Public Economics, vol.3*. North Holland: Elsevier Science Publishers. 4

BEST, MICHAEL CARLOS, & KLEVEN, HENRIK JACOBSEN. 2015. *Housing Market Responses to Transaction Taxes: Evidence From Notches and Stimulus in the UK*. Mimeo: LSE. 4, 28

CAMPBELL, JOHN Y., & MANKIW, N. GREGORY. 1989. The Response of Consumption to Income: A Cross-Country Investigation. *European Economic Review*, **35**(4), 723–756. 4

CHETTY, RAJ, FRIEDMAN, JOHN, OLSEN, TORE, & PISTAFERRI, LUIGI. 2011. Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *Quarterly Journal of Economics*, **126**(2), 749–804. 2, 4, 11

CHODOROW-REICH, GABRIEL. 2014. The Employment Effects of Credit Market Disruptions: Firm-Level Evidence From the 2008-09 Financial Crisis. *Quarterly Journal of Economics*, **129**(1), 1–59. 4

CLOYNE, JAMES, & SURICO, PAOLO. 2013. *Household Debt and the Dynamic Effects of Income Tax Changes*. CEPR Discussion Paper 9649. 21

CLOYNE, JAMES, FERREIRA, CLODOMIRO, & SURICO, PAOLO. 2015. *Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism*. Mimeo: Bank of England. 4

DEFUSCO, ANTHONY, & PACIOREK, ANDREW. 2014. *The Interest Rate Elasticity of Mortgage Demand: Evidence From Bunching at the Conforming Loan Limit*. Mimeo: Wharton School, University of Pennsylvania. 3

DI MAGGIO, MARCO, KERMANI, AMIR, & RAMCHARAN, RODNEY. 2015. *Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging*. Mimeo. 4

EINAV, LIRAN, FINKELSTEIN, AMY, & SCHRIMPF, PAUL. 2015. The Response of Drug Expenditure to Non-Linear Contract Design: Evidence From Medicare Part D. *Quarterly Journal of Economics*, forthcoming. 4

EPSTEIN, LARRY G., & ZIN, STANLEY E. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, **57**(4), 937–969. 30, 47, 53

EVANS, OWEN J. 1983. Tax Policy, the Interest Elasticity of Saving, and Capital Accumulation: Numerical Analysis of Theoretical Models. *American Economic Review*, **73**(3), 398–410. 2, 27

FOLLAIN, JAMES R., & DUNSKY, ROBERT M. 1997. The Demand for Mortgage Debt and the Income Tax. *Journal of Housing Research*, **8**(2), 155–199. 3

GROSS, DAVID B., & SOULELES, NICHOLAS S. 2002. Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence From Credit Card Data. *Quarterly Journal of Economics*, **117**(1), 149–185. 3

GRUBER, JONATHAN. 2006. *A Tax-Based Estimate of the Elasticity of Intertemporal Substitution*. NBER Working Paper 11945. 1, 4

GUVENEN, FATIH. 2006. Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective. *Journal of Monetary Economics*, **53**(7), 1451–1472. 27

HALL, ROBERT E. 1988. Intertemporal Substitution in Consumption. *Journal of Political Economy*, **96**(2), 339–357. 4, 31

JAPPELLI, TULLIO, & PISTAFERRI, LUIGI. 2007. Do People Respond to Tax Incentives? An Analysis of the Italian Reform to the Deductibility of Home Mortgage Interests. *European Economic Review*, **51**, 247–271. 3

JOHNSON, DAVID S., PARKER, JONATHAN A., & SOULELES, NICHOLAS S. 2006. Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, **96**(5), 1589–1610. 4

KEYS, BENJAMIN, PISKORSKI, TOMASZ, SERU, AMIT, & YAO, VINCENT. 2015. *Mortgage Rates, Household Balance Sheets and the Real Economy*. Mimeo. 4

KING, ROBERT G., & REBELO, SERGIO. 1990. Public Policy and Economic Growth: Developing Neoclassical Implications. *Journal of Political Economy*, **98**(5), S126–S150. 4

KLEVEN, HENRIK JACOBSEN, & WASEEM, MAZHAR. 2013. Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *Quarterly Journal of Economics*, **128**(2), 669–723. 2, 3, 4, 8, 11, 12, 14, 15, 16, 18, 19, 25

KOCHERLAKOTA, NARAYANA R. 1996. The Equity Premium: It's Still a Puzzle. *Journal of Economic Literature*, **34**(1), 42–71. 18

KOPCZUK, WOJCIECH, & MUNROE, DAVID. 2014. Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market. *American Economic Journal: Economic Policy*, forthcoming. 12

LAIBSON, DAVID. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, **112**(2), 443–478. 29

MARTINS, NUNO C., & VILLANUEVA, ERNESTO. 2006. The Impact of Mortgage Interest Rate Subsidies on Household Borrowing. *Journal of Public Economics*, **90**, 1601–1623. 3

MEHRA, RAJNISH, & PRESCOTT, EDWARD C. 1985. The Equity Premium: A Puzzle. *Journal of Monetary Economics*, **15**, 145–161. 18

MIAN, ATIF, & SUFI, AMIR. 2009. The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *Quarterly Journal of Economics*, **124**(4), 1449–1496. 4, 31

MIAN, ATIF, & SUFI, AMIR. 2011. House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis. *American Economic Review*, **101**(5), 2132–2156. 4, 18, 28, 31

MIAN, ATIF, & SUFI, AMIR. 2012. The Effects of Fiscal Stimulus: Evidence From the 2009 Cash For Clunkers Program. *Quarterly Journal of Economics*, **127**(3), 1107–1142. 4

MISHKIN, FREDERIC S. 1995. Symposium on the Monetary Transmission Mechanism. *Journal of Economic Perspectives*, **9**(4), 3–10. 4

SAEZ, EMMANUEL. 2010. Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, **2**(3). 4, 18

SIEGEL, JEREMY J., & THALER, RICHARD H. 1997. Anomalies: The Equity Premium Puzzle. *Journal of Economic Perspectives*, **11**(1), 191–200. 18

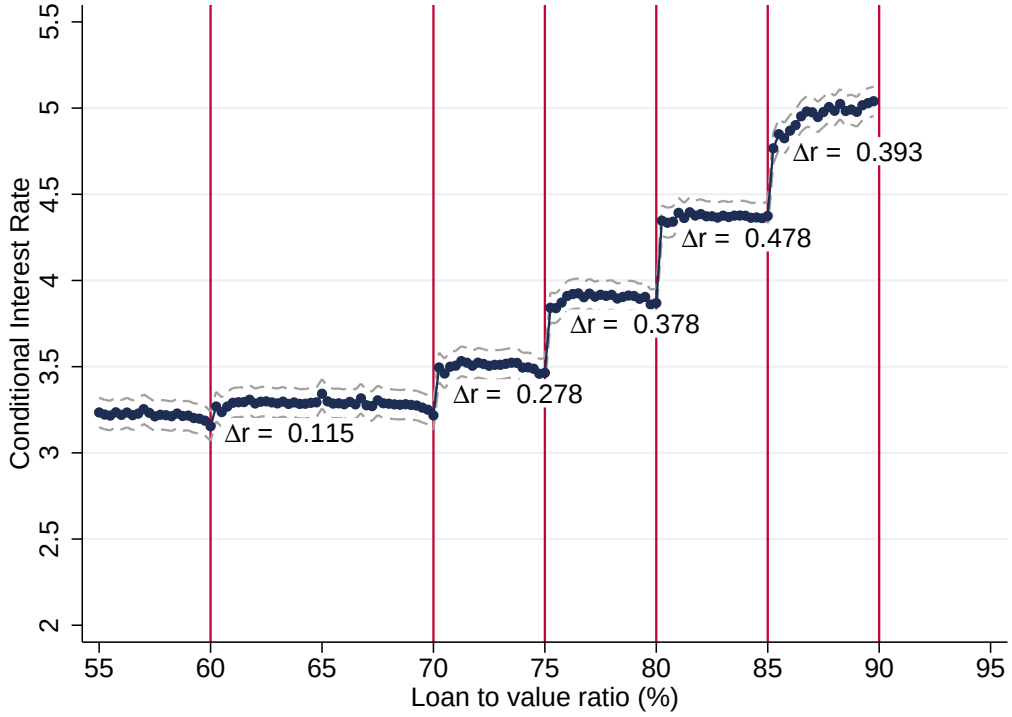
SUMMERS, LAWRENCE H. 1981. Capital Taxation and Accumulation in a Life Cycle Growth Model. *American Economic Review*, **71**(4), 533–544. 2, 4, 27

VISSING-JØRGENSEN, ANNETTE. 2002. Limited Asset Market Participation and the Elasticity of Intertemporal Substitution. *Journal of Political Economy*, **110**(4), 825–853. 4

WEIL, PHILLIPPE. 1990. Nonexpected Utility in Macroeconomics. *Quarterly Journal of Economics*, **105**(1), 29–42. 30, 47, 53

ZELDES, STEPHEN P. 1989. Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, **97**(2), 305–346. 4

FIGURE 1: CONDITIONAL INTEREST RATE JUMPS AT NOTCHES

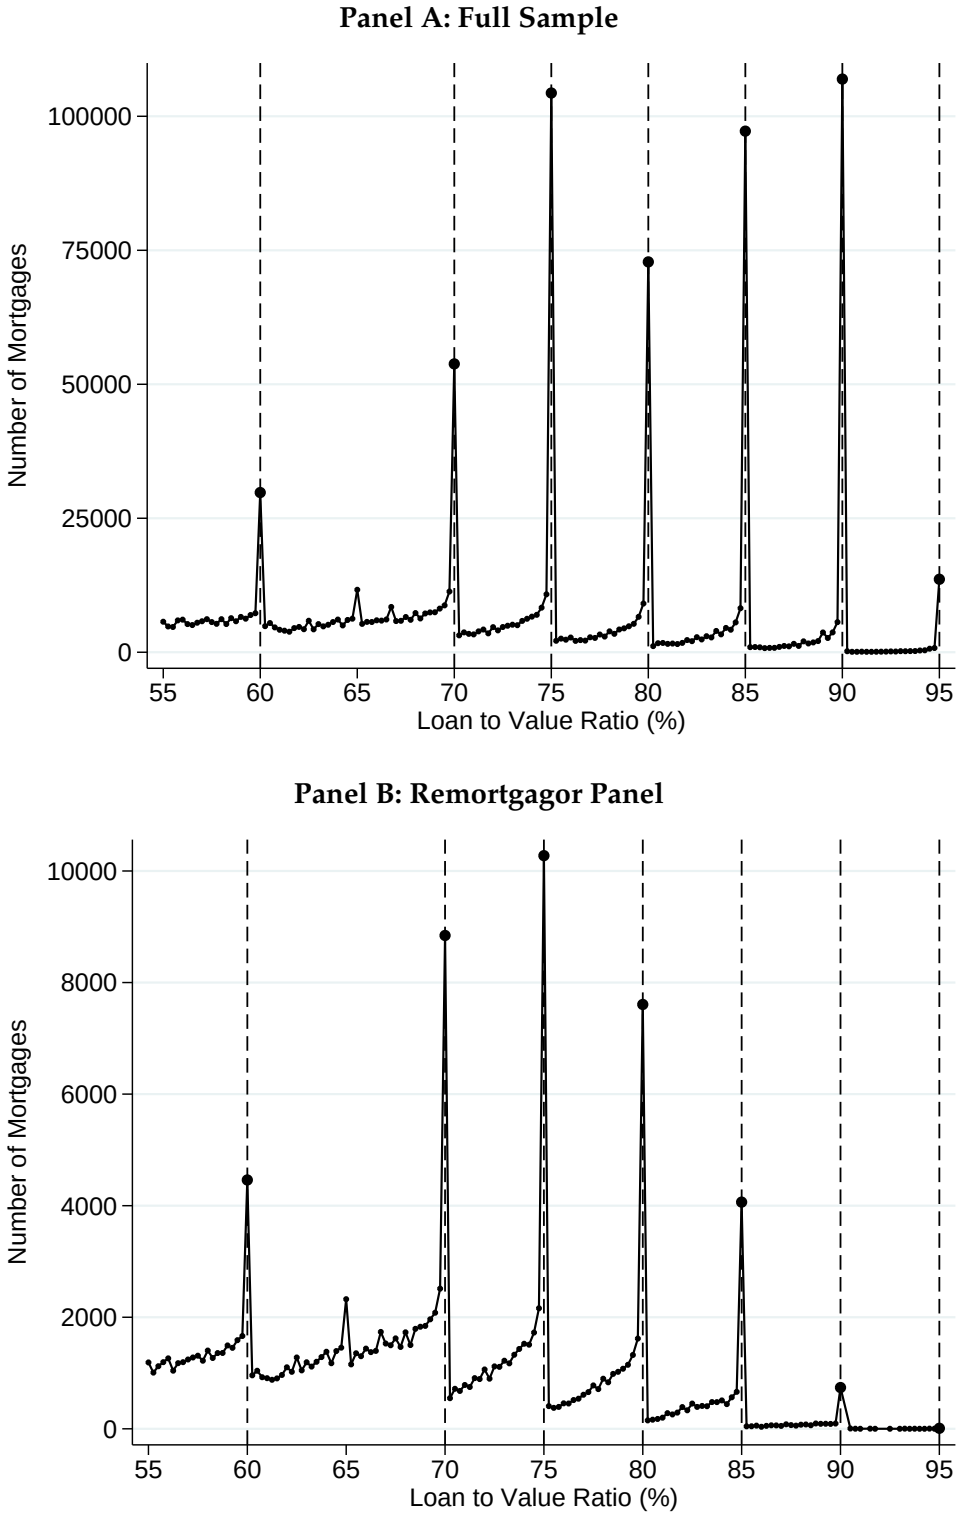


Notes: The figure shows the conditional interest rate as a function of the Loan To Value (LTV) ratio from regression (1):

$$\begin{aligned} r_i \quad = \quad & f\left(LTV_i\right) + \beta_1 \text{lender}_i + \beta_2 \text{type}_i \otimes \text{dur}_i \otimes \text{month}_i + \beta_3 \text{repayment}_i + \beta_4 \text{reason}_i \\ & + s_1\left(\text{age}_i\right) + s_2\left(\text{income}_i\right) I\left\{\text{single}_i\right\} + s_3\left(\text{income}_i\right) I\left\{\text{couple}_i\right\} + s_4\left(\text{term}_i\right) + \nu_i \end{aligned}$$

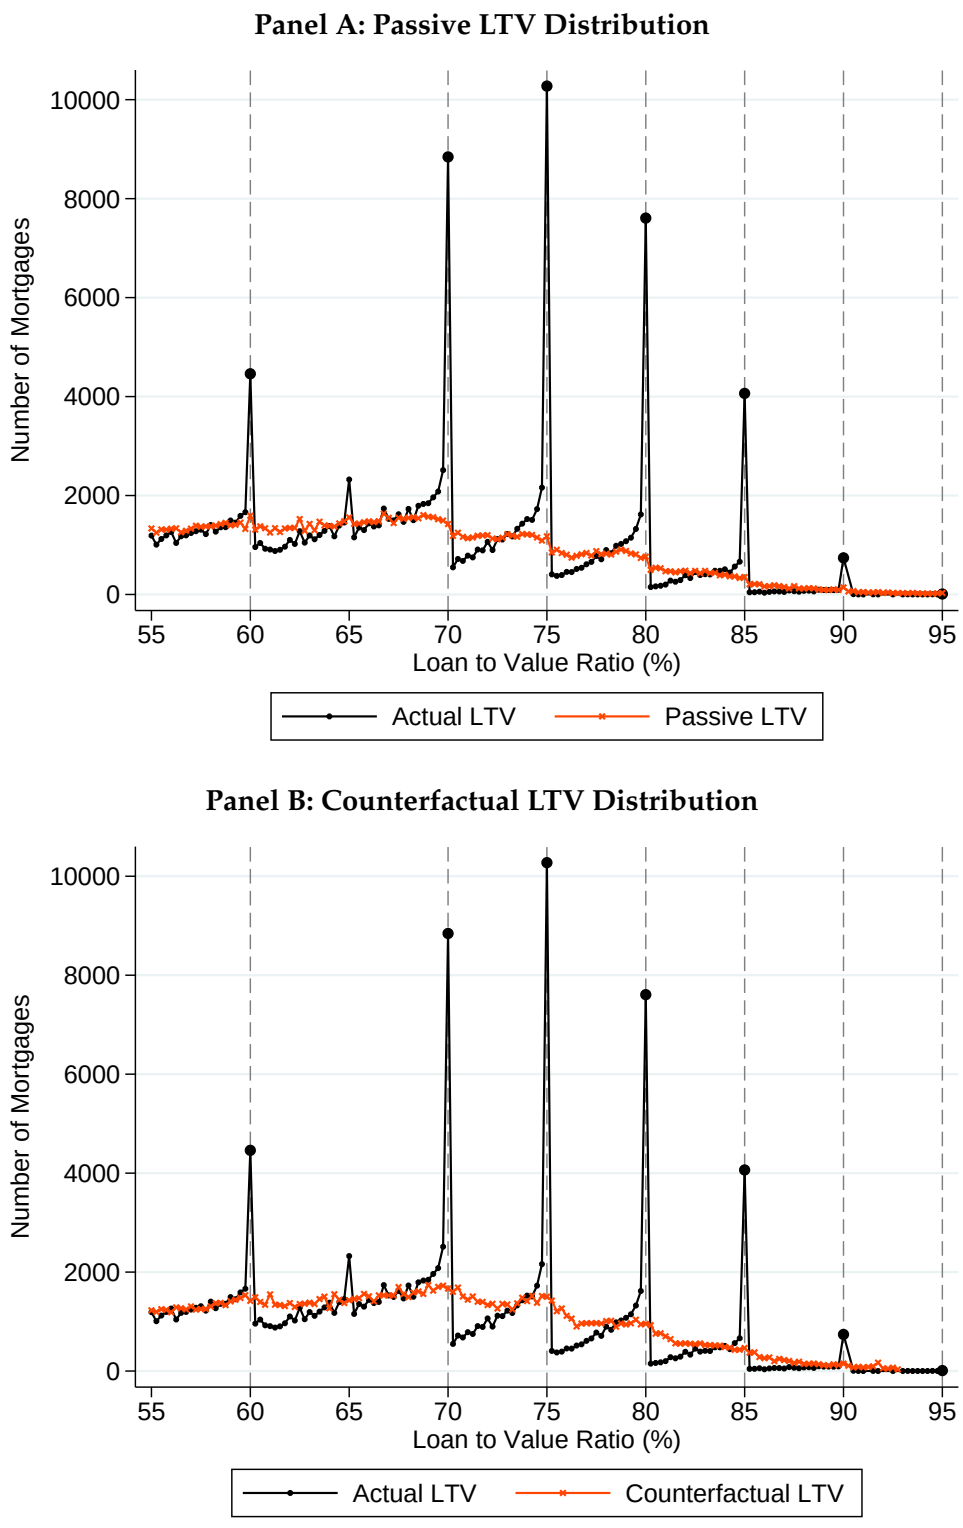
where r_i is the nominal (net) mortgage interest rate for individual i, $f\left(.\right)$ is a step function with steps at each 0.25%-point bin of the LTV, lender_i is a vector of lender dummies, type_i is a vector of mortgage type dummies, dur_i is a vector of dummies for the duration of the initial mortgage period, month_i is a vector of month, repayment_i is a vector of dummies for the repayment type, reason_i is a vector of dummies for the reason for the remortgage, and $s_j\left(.\right)$, $j = 1, ..., 4$ are cubic splines with knots at the quintiles of the distributions of age, income (separately for singles and couples), and the total mortgage term, and $\otimes$ denotes the outer product. In each LTV bin in the figure we plot the coefficient on the LTV bin dummy plus a constant given by the mean predicted value $E\left[\hat{r}_i\right]$ from all the other covariates (i.e., omitting the contribution of the LTV bin dummies). The figure shows that the mortgage interest rate evolves as a step function with sharp jumps at LTV ratios of 60%, 70%, 75%, 80%, and 85%.

FIGURE 2: OBSERVED LTV DISTRIBUTIONS



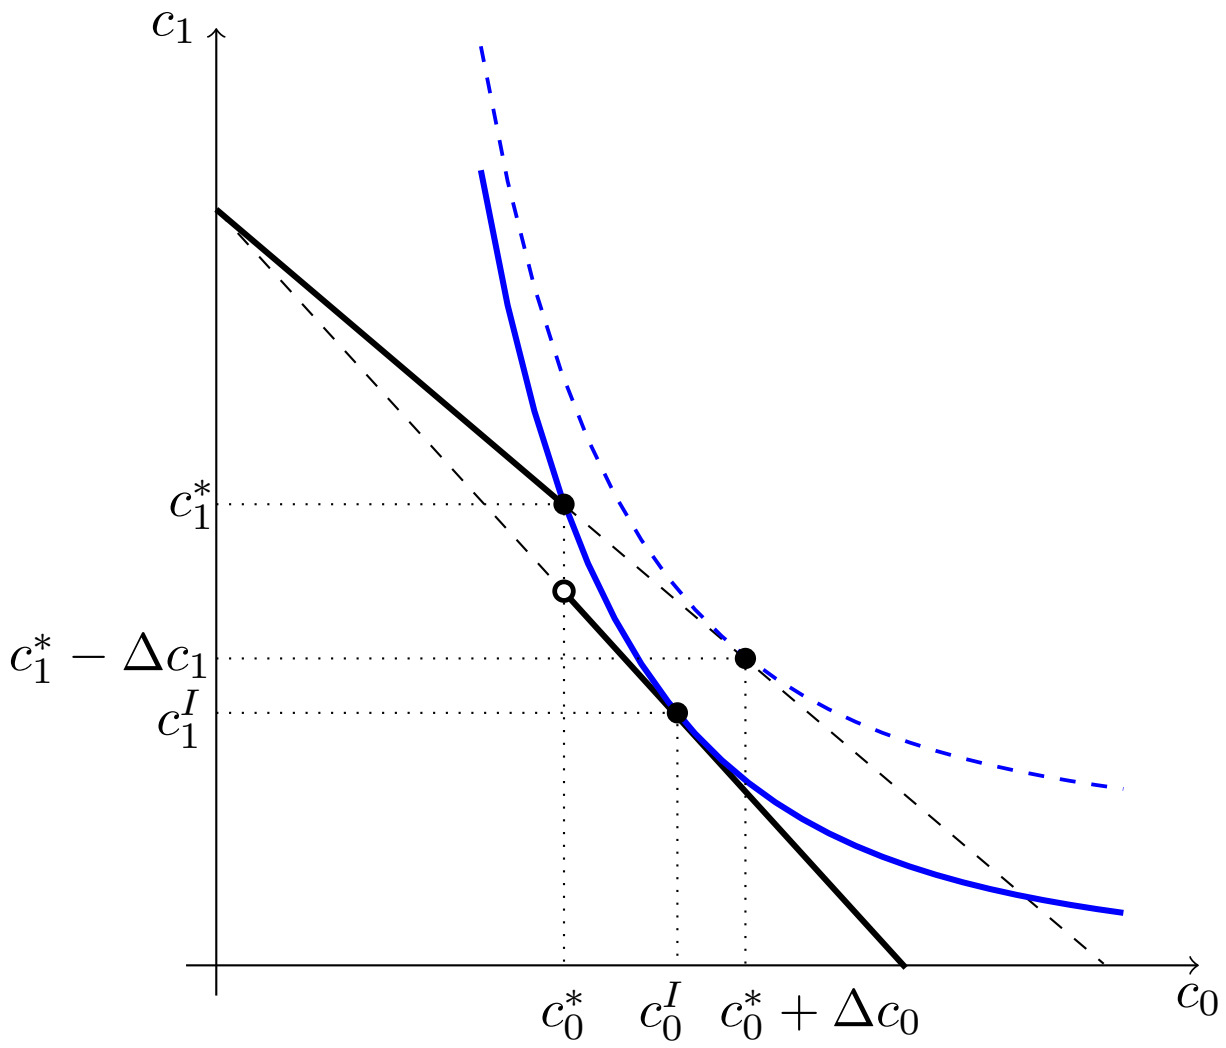
Notes: The figure shows the distribution of Loan to Value (LTV) ratios of mortgages in the full sample of house buyers and remortgagors (panel A) and in the panel of remortgagors (panel B). Black dots denote the number of mortgages in bins 0.25%-points wide, while orange dots denote the number of mortgages in bins with their upper bounds at the notch points (denoted by dashed, vertical lines) at 60%, 70%, 75%, 80%, 85%, 90%, and 95% LTV.

FIGURE 3: CONSTRUCTING THE COUNTERFACTUAL LTV DISTRIBUTION



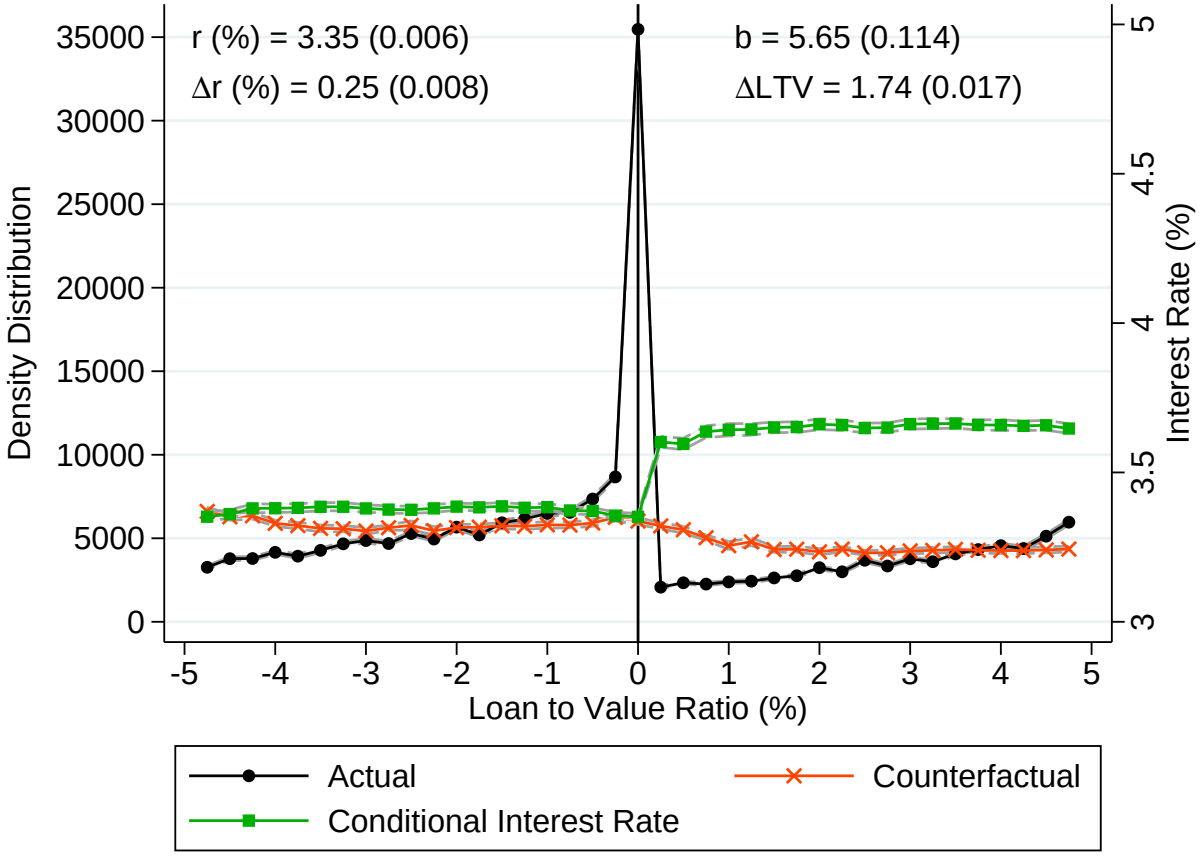
Notes: The figure shows the two steps in the construction of the counterfactual Loan to Value (LTV) distribution in the remortgagor panel. Both panels show the actual LTV distribution in black dots, constructed in the same way as described in the note to figure 2. Panel A shows the distribution of passive LTVs based only on the LTV of the previous mortgage, amortization, and the house value at the time of the remortgage in grey crosses. Panel B shows the distribution of counterfactual LTVs which adjusts passive LTVs for the average equity extraction of non-bunchers in the actual distribution, in grey crosses.

FIGURE 4: THE MARGINAL BUNCHING HOUSEHOLD



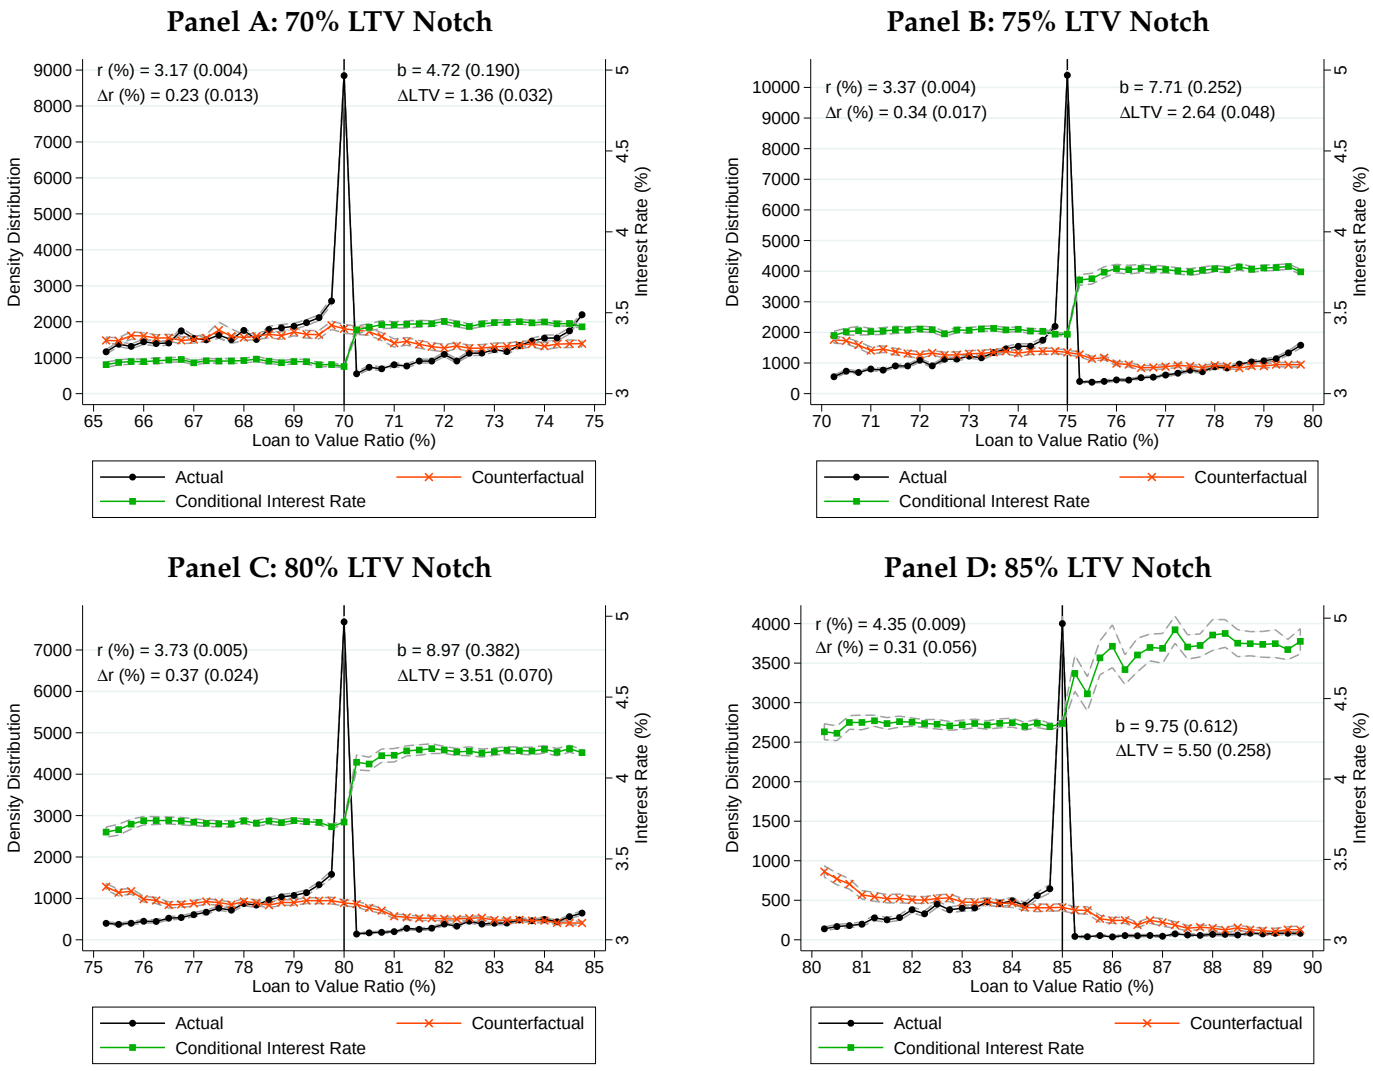
Notes: The figure shows the choice faced by a remortgaging household faced with a notched interest rate schedule in a budget set diagram with consumption in the current period, c_0, on the horizontal axis and consumption in the future, c_1, on the vertical axis. An interest rate notch at a threshold LTV of λ^* implies a threshold level of debt, D^* and current consumption c_0^* above which future consumption drops discretely. The solid black line depicts the budget constraint with slope $-R$ below c_0^* and slope $-(R + \Delta R)$ above c_0^*. The indifference curves shown are those of the marginal bunching household that is indifferent between bunching at the notch and receiving c_0^*, and a point c_0^I in the interior of the higher interest rate bracket. This household would have chosen a point $c_0^* + \Delta c_0$ in the absence of the interest rate notch. As described in sections 3.2 and 5.1, the optimality conditions for c_0^I and $c_0^* + \Delta c_0$ together with the household's indifference between c_0^I and c_0^* identify the curvature of the indifference curves—the Elasticity of Intertemporal Substitution (EIS)—in terms of observable and/or estimable quantities.

FIGURE 5: BUNCHING ESTIMATION POOLING NOTCHES



Notes: The figure shows the actual ($f(\lambda)$ in black dots) and counterfactual ($f_0(\lambda)$ in orange crosses) distributions of LTV (λ), pooling all notches (60%, 70%, 75%, 80%, 85%). The green squares show the conditional interest rate in each LTV bin from regression described in section 2.3 and the footnote to figure 1. The counterfactual LTV distribution is obtained using the method outlined in section 3.1 and the footnote to figure 3. The figure also shows the conditional interest rate for people below the notch, r, and the jump in the conditional interest rate at the notch Δr on the left. On the right, the figure shows the normalized amount of bunching in the actual distribution b, calculated as described in section 3.2, as $\int_{\lambda_-}^{\lambda^*} [f(\lambda) - f_0(\lambda)]\, d\lambda / f_0(\lambda^*)$ where λ^* is the position of the notch, which is normalized to 0 in this figure, and λ_- is the lower bound of the excluded range, for which we use -1 in our estimation.

FIGURE 6: BUNCHING AT INDIVIDUAL NOTCHES



Notes: The figure shows the actual ($f(\lambda)$ in black dots) and counterfactual ($f_0(\lambda)$ in orange crosses) distributions of LTV (λ), at the notches at 70% (Panel A), 75% (Panel B), 80% (Panel C) and 85% (Panel D) LTV. The green squares show the conditional interest rate in each LTV bin from regression described in section 2.3 and the footnote to figure 1. The counterfactual LTV distribution is obtained using the method outlined in section 3.1 and the footnote to figure 3. The figure also shows the conditional interest rate for people below the notch, r, and the jump in the conditional interest rate at the notch Δr on the left. On the right, the figure shows the normalized amount of bunching in the actual distribution b, calculated as described in section 3.2, as $\int_{\lambda_-}^{\lambda^*} [f(\lambda) - f_0(\lambda)]\, d\lambda / f_0(\lambda^*)$ where λ^* is the position of the notch, which is normalized to 0 in this figure, and λ_- is the lower bound of the excluded range, for which we use -1 in our estimation.

TABLE 1: SUMMARY STATISTICS

	(1) **Full sample** mean/sd	(2) **Remortgagors** mean/sd	(3) **Remortgagor Panel** mean/sd
Interest Rate (%)	4.11	4.25	3.99
	(1.31)	(1.39)	(1.24)
Loan Size (£)	144,598.0	140,563.2	149,144.9
	(121,464.9)	(121,455.5)	(117,560.8)
Property Value (£)	260,160.7	265,574.6	268,813.9
	(256,564.6)	(270,117.5)	(237,418.6)
Loan to Value Ratio (%)	60.0	56.9	58.8
	(21.9)	(21.4)	(19.7)
Gross Income (£)	57,718.5	57,975.8	58,702.0
	(83,160.2)	(83,960.3)	(76,309.9)
Loan to Income Ratios	2.79	2.71	2.78
	(1.81)	(2.00)	(1.29)
Repayments to Income (%)	21.9	22.7	21.7
	(40.5)	(49.5)	(12.5)
Proportion with Joint Income	0.54	0.55	0.55
	(0.50)	(0.50)	(0.50)
Mortgage Term (years)	20.7	19.1	20.1
	(7.37)	(6.87)	(6.63)
Proportion of Fixed Rate Mortgages	0.68	0.65	0.72
	(0.47)	(0.48)	(0.45)
Proportion of Remortgage Products	0.57	0.88	0.76
	(0.49)	(0.32)	(0.43)
Borrower's Age	40.0	41.6	40.2
	(9.95)	(9.53)	(8.80)
Observations	2,841,309	1,854,430	564,882

Notes: The table shows summary statistics of the three samples used in our analysis. Column (1) shows our full sample: all usable mortgages in the PSD (house purchases and remortgages) including (a) observations where we can find information on the product fee in MoneyFacts (b) any remortgagor's previous mortgage (which we can use without fee information). Column (2) restricts the sample to remortgagors only (all remortgage products and any remortgagor's previous mortgage). Finally, column (3) shows the sample of remortgagors in the panel (remortgagors we can link over time and where we can construt a counterfactual LTV ratio).

TABLE 2: FROM BUNCHING TO MORTGAGE DEMAND ELASTICITIES

Statistic	Notch					
	60	70	75	80	85	Pooled
Panel A: Raw Bunching Evidence						
$r(\%)$	3.08	3.17	3.37	3.73	4.35	3.35
	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
$\Delta r(\%)$	0.12	0.23	0.34	0.37	0.31	0.25
	(0.01)	(0.01)	(0.02)	(0.02)	(0.06)	(0.01)
b	2.00	4.72	7.71	8.97	9.75	5.65
	(0.13)	(0.19)	(0.25)	(0.38)	(0.61)	(0.11)
ΔLTV	0.54	1.36	2.64	3.51	5.50	1.74
	(0.03)	(0.03)	(0.05)	(0.07)	(0.26)	(0.02)
Panel B: Descriptive Statistics						
L^*	178,538.25	181,863.03	168,813.06	160,525.39	146,879.70	170,048.94
	(1,333.37)	(1,032.00)	(734.82)	(710.09)	(991.63)	(463.59)
V^*	298,822.56	260,326.70	225,294.78	200,861.78	172,969.72	236,388.34
	(2,231.82)	(1,476.34)	(983.02)	(887.24)	(1,167.37)	(661.44)
Δt (yrs)	3.27	3.21	3.24	3.56	3.65	3.34
	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)
ΔEquity$/V$ (%)	3.34	4.11	5.26	6.12	5.89	4.88
	(0.09)	(0.06)	(0.07)	(0.07)	(0.12)	(0.04)
Panel C: Mortgage Demand Elasticities						
Δc_0	1,611	3,543	5,946	7,050	9,506	4,104
	(87.72)	(83.22)	(113.97)	(145.09)	(455.76)	(39.86)
Δc_1	2,418	5,065	8,181	9,807	12,241	5,802
	(105.98)	(105.92)	(150.37)	(237.88)	(677.98)	(66.36)
r^* (%)	13.20	11.78	10.35	9.71	7.18	10.92
	(1.11)	(0.62)	(0.46)	(0.47)	(0.81)	(0.27)
Elasticity ε	0.07	0.19	0.40	0.56	1.37	0.25
	(0.01)	(0.01)	(0.02)	(0.03)	(0.26)	(0.01)

Notes: The table shows our reduced form estimates using bunching at the various LTV notches separately, and pooling the notches from 60% to 85%. r is the conditional nominal interest rate below the notch, Δr is the interest rate jump at the notch, estimated as described in section 2.3, b is our normalized bunching estimate, and ΔLTV is our estimate of the leverage response, estimated as described in section 3.2. L^* and V^* are the average loan size and property value among individuals less than 1% below the notch, respectively. Δt is the average time between mortgages (in years), and ΔEquity $/ V$ is the average equity extracted as a proportion of the house value. Δc_0 is the average foregone current consumption by bunchers, and Δc_1 is the average increased consumption in the future for bunchers, implying an annualized rate of return of $r^* = (\Delta c_1/\Delta c_0)^{1/\Delta t} - 1$. Together, r, r^* and ΔLTV imply a reduced-form mortgage demand elasticity of $\varepsilon = (\Delta\text{LTV}/\text{LTV}^*) \,/\, [(r^* - r) \,/\, (1 + r)]$. The standard errors, shown in parentheses, are obtained by bootstrapping the estimation routine, stratifying by notch, 100 times.

TABLE 3: FROM BUNCHING TO ELASTICITIES OF INTERTEMPORAL SUBSTITUTION

Statistic	Notch					
	60	**70**	**75**	**80**	**85**	**Pooled**
b	2.00	4.72	7.71	8.97	9.75	5.65
	(0.13)	(0.19)	(0.25)	(0.38)	(0.61)	(0.11)
a	0.67	0.17	0.33	0.14	0.05	0.36
	(0.11)	(0.05)	(0.06)	(0.03)	(0.02)	(0.04)
b_{Adj}	6.12	5.71	11.53	10.39	10.27	8.76
	(1.61)	(0.44)	(1.06)	(0.63)	(0.65)	(0.65)
$\Delta\lambda$	1.68	1.70	4.08	4.22	6.15	2.85
	(0.45)	(0.15)	(0.38)	(0.23)	(0.34)	(0.22)
EIS σ	0.10	0.05	0.15	0.11	0.23	0.08
	(0.05)	(0.01)	(0.03)	(0.01)	(0.04)	(0.01)

Notes: The table shows the results of our structural estimation as described in section 5.2. b is our normalized bunching estimate as described in section 3.2 and the footnote to figure 6. a is the adjustment factor for optimization frictions (the number of individuals observed in the dominated region defined by $\sigma = 0$ in the actual distribution divided by the number of individuals in the same region in the counterfactual distribution), and $b_{Adj} = b/(1-a)$ is our bunching estimate, adjusted for optimization frictions. $\Delta\lambda$ is the leverage response corresponding to b_{Adj}. σ is the Elasticity of Intertemporal Substitution (EIS) that solves the marginal buncher's indifference condition (20). The standard errors, shown in parentheses, are obtained by bootstrapping the estimation routine, stratifying by notch, 100 times.

TABLE 4: HETEROGENEITY IN THE ELASTICITY OF INTERTEMPORAL SUBSTITUTION

Covariate	Quartile			
	1	**2**	**3**	**4**
Age	0.05	0.06	0.10	0.09
	(0.01)	(0.02)	(0.04)	(0.03)
Income	0.09	0.08	0.05	0.07
	(0.03)	(0.02)	(0.01)	(0.04)
Loan to Income	0.05	0.06	0.06	0.10
	(0.03)	(0.02)	(0.01)	(0.03)
House Price Growth Rate	0.07	0.05	0.04	0.24
	(0.03)	(0.01)	(0.01)	(0.06)
Interest Rate Change (Passive)	0.04	0.05	0.08	0.11
	(0.01)	(0.01)	(0.02)	(0.05)
Interactions With Age:				
Age $\times$ Income	0.08	0.06	0.05	0.14
	(0.03)	(0.02)	(0.02)	(0.04)
Age $\times$ Loan to Income	0.02	0.08	0.07	0.21
	(0.01)	(0.02)	(0.01)	(0.05)
Age $\times$ House Price Growth	0.11	0.04	0.05	0.26
	(0.03)	(0.01)	(0.01)	(0.06)
Age $\times$ Interest Rate Change	0.07	0.05	0.07	0.13
	(0.04)	(0.01)	(0.01)	(0.03)

Notes: The table shows the heterogeneity in our estimated EIS σ (using the pooled average notch) by age, income, loan to income (LTI), house price growth, interest rate change since the previous mortgage (assuming passive borrower behavior). The table also shows heterogeneity for each of these covariates interacted with age. For each covariate, we partition the remortgagor panel into 4 quartiles and separately estimate σ in each quartile. The standard errors, shown in parentheses, are obtained by bootstrapping the estimation routine, stratifying by notch, 100 times.

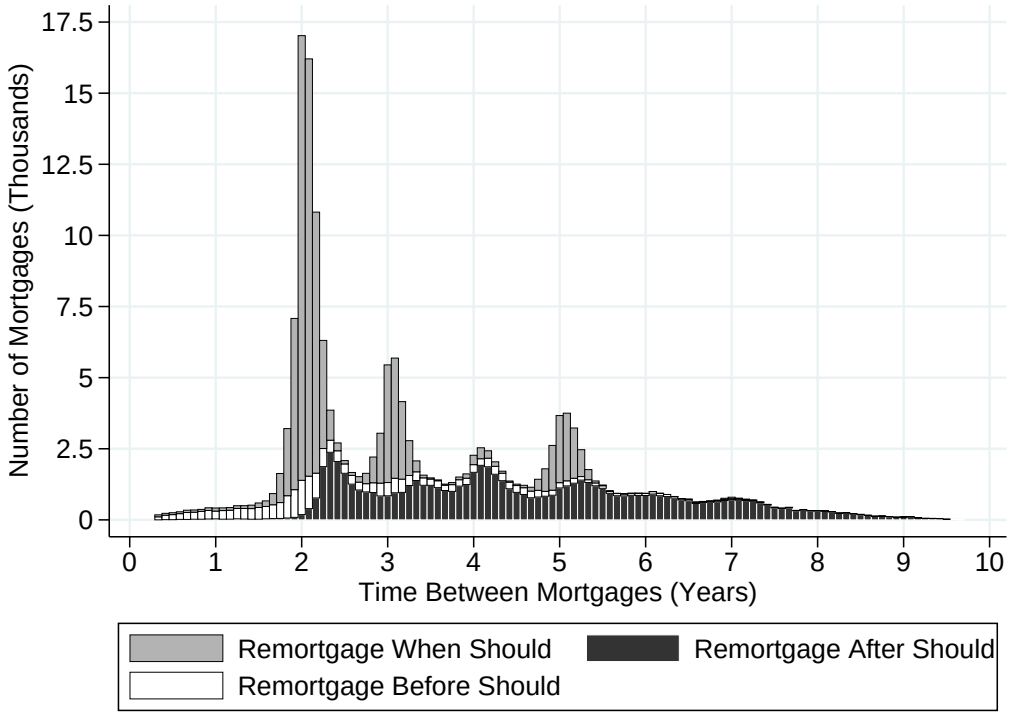
TABLE 5: ROBUSTNESS OF ELASTICITY OF INTERTEMPORAL SUBSTITUTION ESTIMATES

Panel A: Future Interest Rates $\{r_t\}_2^\infty$				
	2%	$r_1 = 3.35\%$	7%	15%
σ	0.0743	0.0764	0.0817	0.0909
	(0.0111)	(0.0115)	(0.0123)	(0.0138)
$\tilde{R}$	79.0507	18.8395	6.6208	3.1167
	(10.8433)	(0.4790)	(0.0134)	(0.0228)
Panel B: Discount Factor δ				
	0.7	0.9	0.95	0.99
σ	0.0811	0.0772	0.0764	0.0759
	(0.0133)	(0.0118)	(0.0115)	(0.0113)
$\tilde{R}$	7.7614	15.0856	18.8395	23.2255
	(0.7344)	(0.7306)	(0.4790)	(0.1150)
Panel C: Hyperbolic Discounting β				
	0.5	0.7	0.9	1
σ	0.0766	0.0765	0.0765	0.0764
	(0.0115)	(0.0115)	(0.0115)	(0.0115)
$\tilde{R}$	18.8328	18.8361	18.8385	18.8395
	(0.4812)	(0.4801)	(0.4793)	(0.4790)
Panel D: Risk Aversion γ (With Uncertainty)				
	0	5	10	100
σ	0.0764	0.0795	0.0781	0.0749
	(0.0115)	(0.0120)	(0.0117)	(0.0112)
$\tilde{R}$	18.8395	9.0514	11.8217	39.4246
	(0.4790)	(0.0524)	(0.1340)	(2.5203)

Notes: The table shows the robustness of our estimates of the elasticity of intertemporal substitution σ to a number of the assumptions of our structural model. Each panel shows our estimate of σ for a different parameter assumption and the effect of changing the assumption on $\tilde{R} = 1/\left(1 - \beta^\sigma R^{\sigma-1}\right)$ together with bootstrapped standard errors. Panel A relaxes the assumption that future (nominal) interest rates $\{r_t\}_2^\infty$ are equal to the current rate r_1 by considering future nominal (our estimates incorporate 2% inflation) interest rates from 2% to 15%. Panel B considers discount factors ranging from 0.7, through our baseline assumption of 0.95, to 0.99. Panel C relaxes our assumption that households discount the future with traditional geometric discounting, and allows for quasi-hyperbolic $\beta - \delta$ discounting with present bias parameters β from 0.5 to 1 (standard discounting). Panel D relaxes our assumption of CES/CRRA utility and allows for separate parameters governing risk aversion (γ) and intertemporal substitution (σ) as in Epstein & Zin (1989) and Weil (1990) (see appendix C for detailed derivations) exploring risk aversion parameters from 0 (risk neutrality) to 100 (extreme risk aversion). In all cases, while the parameters considered have large effects on $\tilde{R}$, the estimates of σ are remarkably robust.

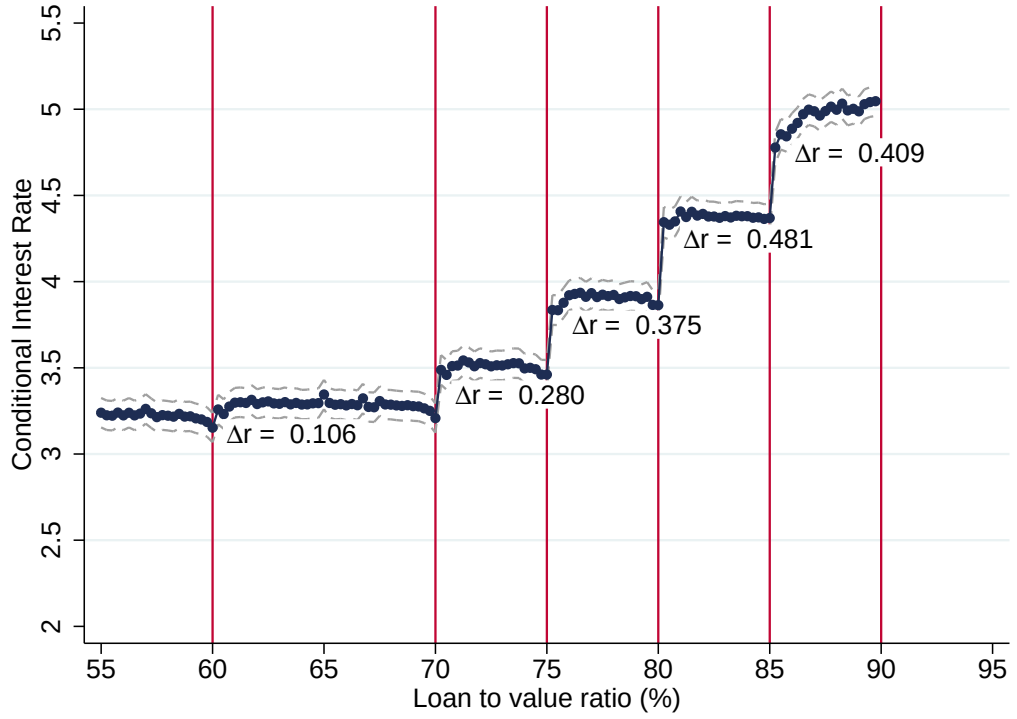
A Appendix Figures & Tables

FIGURE A.1: REMORTGAGES HAPPEN WHEN THE RESET RATE KICKS IN



Notes: The figure shows the distribution of the time to remortgage, excluding individuals where the date on which the reset rate kicks in is unobserved. The figure shows individuals who remortgage within 60 days of their reset rate kicking in in green, individuals who remortgage more than 60 days after their reset rate kicks in in blue, and individuals who remortgage more than 60 days before their reset rate kicks in in orange.

FIGURE A.2: CONDITIONAL INTEREST RATE JUMPS WITHOUT BORROWER DEMOGRAPHICS

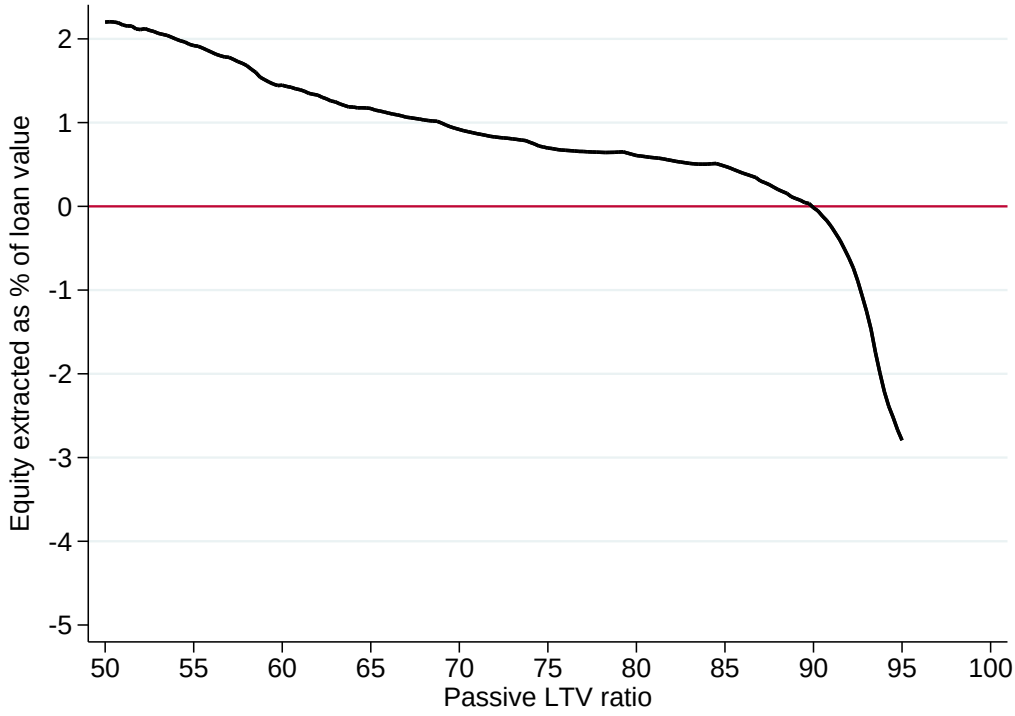


Notes: The figure shows the conditional interest rate as a function of the Loan To Value (LTV) ratio from a regression like (1) as shown in figure 1, but omitting borrower characteristics:

$$r_i \quad = \quad f\left(LTV_i\right) + \beta_1 \text{lender}_i + \beta_2 \text{type}_i \otimes \text{dur}_i \otimes \text{month}_i + \beta_3 \text{repayment}_i + \beta_4 \text{reason}_i + s_4\left(\text{term}_i\right) + \nu_i$$

where r_i is the nominal (net) mortgage interest rate for individual i, $f\left(.\right)$ is a step function with steps at each 0.25%-point bin of the LTV, lender_i is a vector of lender dummies, type_i is a vector of mortgage type dummies, dur_i is a vector of dummies for the duration of the initial mortgage period, month_i is a vector of month, repayment_i is a vector of dummies for the repayment type, reason_i is a vector of dummies for the reason for the remortgage, and $s_4\left(.\right)$,is a cubic spline with knots at the quintiles of the distribution of the total mortgage term, and $\otimes$ denotes the outer product. In each LTV bin in the figure we plot the coefficient on the LTV bin dummy plus a constant given by the mean predicted value $E\left[\hat{r}_i\right]$ from all the other covariates (i.e., omitting the contribution of the LTV bin dummies). The figure shows that the mortgage interest rate evolves as a step function with sharp jumps at LTV ratios of 60%, 70%, 75%, 80%, and 85%.

FIGURE A.3: EQUITY EXTRACTION BY PASSIVE LTV



Notes: The figure shows the moving average of equity extraction on the y axis, calculated among households that do not bunch in the actual LTV distribution. The x axis is the *passive LTV*, the LTV that results from applying the amortization to the previous mortgage and using the new lender-assessed property valuation. This moving average is used to adjust the passive LTV distribution to obtain the counterfactual LTV distribution.

TABLE A.1: PARAMETERS FOR STRUCTURAL ESTIMATION

Parameter	Notch					
	60	**70**	**75**	**80**	**85**	**Pooled**
y_0	48,894.10	46,929.45	43,202.83	41,468.96	39,659.52	44,413.05
	(492.43)	(397.96)	(231.68)	(182.29)	(237.44)	(162.80)
$R_0 D_0$	168,970.20	171,344.61	156,860.92	147,657.06	136,031.48	158,783.11
	$(1,348.37)$	(971.79)	(709.30)	(621.31)	(890.40)	(450.05)
D_1^*	178,538.25	181,863.03	168,813.06	160,525.39	146,879.70	170,048.94
	$(1,333.37)$	$(1,032.00)$	(734.82)	(710.09)	(991.63)	(463.59)
$\{r_t\}_2^\infty$ (%)	3.08	3.17	3.37	3.73	4.35	3.35
	(0.01)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
δ	0.95	0.95	0.95	0.95	0.95	0.95
Δt (yrs)	3.27	3.21	3.24	3.56	3.65	3.34
	(0.02)	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)

Notes: The table shows the parameters used in the structural estimation of the Elasticity of Intertemporal Substitution (EIS), the contents of X in the indifference equation (20). All parameters are averages of the individuals in the remortgagor panel whose LTV is less than 1% below a notch, trimming the top and bottom 5% of equity extracted to minimize the influence of outlier equity extraction. y_0 is average income, $R_0 D_0$ is the average outstanding balance, D_1^* is the average loan size, δ is the households' discount factor, and Δt is the average time between mortgages.

B Dominated Region

The upper bound of the dominated region is given by the $\Delta\lambda$ that solves the indifference condition for the marginal buncher (20) when $\sigma = 0$. Combining equations (14) and (13), consumption utility at the marginal buncher's most preferred interior point can be written as

$$C^I = D_1^* \left[\frac{D_1^{CF}}{D_1^*} \left(\frac{R}{\tilde{R}} + (\beta R)^\sigma \right) + \left(\frac{R+\Delta R}{R} + (\beta R)^\sigma \right) \frac{y_0 - R_0 D_0}{D_1^*} \right] \times \left[\left(\frac{R+\Delta R}{\tilde{R}} \right)^{1-\sigma} + \left(\beta \tilde{R}\right)^\sigma \right]^{\frac{1}{\sigma-1}}$$

Similarly, using (17) and (18), consumption utility at the notch is

$$C^N = D_1^* \left[\left(1 + \frac{y_0 - R_0 D_0}{D_1^*} \right)^{\frac{\sigma-1}{\sigma}} + \beta \tilde{R} \left(\frac{D_1^{CF}}{D_1^*} \left(\frac{R}{\tilde{R}} + (\beta R)^\sigma \right) + (\beta R)^\sigma \frac{y_0 - R_0 D_0}{D_1^*} - \frac{R}{\tilde{R}} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

Combining these, and noting that $D_1^{CF}/D_1^* = 1 + \Delta\lambda/\lambda^*$, for general $\sigma \geq 0$, the indifference condition $C^I = C^N$ is given by

$$\begin{aligned} & \left[\left((\beta R)^\sigma + \frac{R}{\tilde{R}} \right) \left(1 + \frac{y_0 - R_0 D_0}{D_1^*} + \frac{\Delta\lambda}{\lambda^*} \right) + \frac{\Delta R}{\tilde{R}} \frac{y_0 - R_0 D_0}{D_1^*} \right] \times \left[\left(\frac{R+\Delta R}{\tilde{R}} \right)^{1-\sigma} + \left(\beta \tilde{R}\right)^\sigma \right]^{\frac{1}{\sigma-1}} \\ = & \left[\left(1 + \frac{y_0 - R_0 D_0}{D_1^*} \right)^{\frac{\sigma-1}{\sigma}} + \beta \tilde{R} \left((\beta R)^\sigma \left(1 + \frac{y_0 - R_0 D_0}{D_1^*} + \frac{\Delta\lambda}{\lambda^*} \right) + \frac{R}{\tilde{R}} \frac{\Delta\lambda}{\lambda^*} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \end{aligned}$$

As $\sigma \to 0$, the right hand side of this expression becomes a minimum function, and so the indifference condition becomes

$$\begin{aligned} \frac{\left(1 + \frac{R}{\tilde{R}}\right) \left(1 + \frac{y_0 - R_0 D_0}{D_1^*} + \frac{\Delta\lambda}{\lambda^*} \right) + \frac{\Delta R}{\tilde{R}} \frac{y_0 - R_0 D_0}{D_1^*}}{1 + \frac{R+\Delta R}{\tilde{R}}} &= \min \left\{ 1 + \frac{y_0 - R_0 D_0}{D_1^*}, 1 + \frac{y_0 - R_0 D_0}{D_1^*} + \frac{\Delta\lambda}{\lambda^*} \left(1 + \frac{R}{\tilde{R}} \right) \right\} \\ &= 1 + \frac{y_0 - R_0 D_0}{D_1^*} \end{aligned}$$

where the second equality follows since $\Delta\lambda$, λ^*, R, and $\tilde{R}$ are all positive. Solving this yields

$$\Delta\lambda = \lambda^* \frac{\Delta R}{R + \tilde{R}} \tag{B.1}$$

as the upper bound of the dominated region.

C Uncertainty and Epstein-Zin-Weil Preferences

Consider households identical to the remortgagors we have studied thus far, but with preferences over consumption defined recursively as in Epstein & Zin (1989) and Weil (1990), given in equation (21). We assume that households know their future incomes $\left\{y_t^i\right\}_{t=0}^{\infty}$ with certainty, but that future interest rates are stochastic. At time zero, household i faces a mortgage interest notch, as described in Section 5.1. The household can borrow at a rate of $R_1^i = R$ if it chooses an LTV of $\lambda \leq \lambda^*$and at a rate of $R_1^i = R + \Delta R$ if it chooses $\lambda > \lambda^*$. In all future periods $t > 0$, the household refinances at an interest rate that is drawn from an i.i.d. distribution and that is revealed at the beginning of the period. The assumption of i.i.d. distributed interest rates and deterministic income can be readily relaxed in computational analysis, but allow for an analytical expression of the estimating indifference equation and are therefore expositionally convenient.

Optimization in periods $t \geq 1$: As in the benchmark model, we begin by solving the maximization problem in future periods $t \geq 1$. Household i enters period t with debt $R_t^i D_t^i$ and income of y_t^i. It observes the realization of the mortgage interest rate R_{t+1}^i and chooses consumption c_t^i and debt D_{t+1}^i to maximize (21) subject to the budget constraint (5). It is convenient to denote the household's human wealth—the expected net present value of its lifetime earnings—as

$$Y_t^i \equiv y_t + E_t \left\{ \sum_{s=t+1}^{\infty} y_s^i \prod_{j=t+1}^{s} \left(R_j^i\right)^{-1} \right\}.$$

This is a stochastic version of the term Y^i in the benchmark model. It can be shown that the optimal choice of consumption (full derivation available upon request) is to consume $c_t = \mu_t \left(Y_t^i - R_t^i D_t^i\right)$, where μ_t is the marginal propensity to consume out of total (financial and human) wealth, given by

$$\mu_t = \frac{1 - \delta^\sigma \hat{R}^{\sigma-1}}{1 + \delta^\sigma \left(R_{t+1}^{\sigma-1} - \hat{R}^{\sigma-1}\right)}, \tag{C.2}$$

and $\hat{R} \equiv \left(E_{t-1}\left\{R_{t+1}^{1-\gamma}\right\}\right)^{\frac{1}{1-\gamma}}$ is the certainty-equivalent expected interest rate.

The standard result, that the marginal propensity to consume out of total wealth is constant, is slightly modified in our case by the fact that the borrowing rate is revealed at the beginning of the period. The time-dependence of the marginal propensity to consume is captured by the $R_{t+1}^{\sigma-1} - \hat{R}^{\sigma-1}$ term in the denominator that implies that households consume more when interest rates are unexpectedly low. If interest rates were known with certainty, then $R_{t+1}^{\sigma-1} = \hat{R}^{\sigma-1}$, and

$\mu_t = \tilde{R}^{-1}$. Without risk aversion ($\gamma = 0$) consumption in periods $t \geq 1$ would correspond exactly to that in the benchmark model, as in (10).

Using the optimal choice of consumption and the budget constraint (5), we can approximate the continuation value C_t in (21) as

$$C_t = \frac{\left((1-\delta)\,\mu_t^{-\frac{1}{\sigma}}\right)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}}\left(Y_t^i - R_t^i D_t^i\right)^{1-\gamma} - 1}{(1-\delta)\,(1-\gamma)}. \tag{C.3}$$

We will use the continuation value C_t in evaluating choices in time zero. The optimal choice of consumption at time zero maximizes (21) for $t = 0$, subject to the relevant constraint (with or without a notch), where $C_1$is given by (C.3).

OPTIMIZATION AT $t = 0$ WITHOUT A NOTCH: We begin by studying optimization without a notch, i.e. the household can borrow any amount at an interest rate $R_1^i = R$ at time zero. For simplicity (but without loss of generality) assume that the time zero interest rate in the counterfactual without a notch is the certainty equivalent interest rate $R = \hat{R}$. Optimal consumption is given in the counterfactual without a notch is by

$$c_0^{CF} = \mu_0\left(Y_0 - R_0 D_0\right), \tag{C.4}$$

where μ_0 is the marginal propensity to consume, given by (C.2), so that

$$c_0^{CF} = \left(1 - \beta^\sigma \hat{R}^{\sigma-1}\right)\left(Y_0 - R_0 D_0\right).$$

We surpress the household index i for notational ease.
Using the budget constraint (5) we obtain that the counterfactual debt choice is given by

$$D_1^{CF} = D^* + \Delta D = \delta^\sigma \hat{R}^{\sigma-1} R_0 D_0 + \left(1 - \delta^\sigma \hat{R}^{\sigma-1}\right) Y_0 - y_0. \tag{C.5}$$

OPTIMIZATION AT $t = 0$ WITH A NOTCH: We consider an indifferent household–one that obtains equal utility when bunching at the notch λ^* and when choosing the optimal interior LTV choice. If the indifferent household chooses not to bunch, it will chose an interior choice of consumption at the higher interest rate of $R + \Delta R$. As before, the optimal interior choice of consumption is to

consume a share μ_0 of total wealth, given by (C.2). This gives

$$c_0^I = \frac{1-\delta^\sigma \hat{R}^{\sigma-1}}{1+\delta^\sigma\left((R+\Delta R)^{\sigma-1}-\hat{R}^{\sigma-1}\right)}\left(Y_0 - R_0 D_0\right).$$

As in the benchmark model in Section 5.1, we can use the counterfactual choice of debt (C.5) to eliminate human wealth Y_0 from this equation:

$$c_0^I = \frac{D^* + \Delta D - (R_0 D_0 - y_0)}{1+\delta^\sigma\left((R+\Delta R)^{\sigma-1}-\hat{R}^{\sigma-1}\right)}, \tag{C.6}$$

Using the value function ((21)), its functional form (C.3), and formula for the marginal propensity to consume (C.2) , we can write

$$(1-\delta)(1-\gamma)C_0^I \cong (1-\delta)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}}\left(c_0^I\right)^{1-\gamma}\left(1+\frac{\delta^\sigma(R+\Delta R)^{\sigma-1}}{1-\delta^\sigma\hat{R}^{\sigma-1}}\right)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}} - 1$$

If, in contrast, the household chooses to bunch at the notch, giving $D_1^N = D^* \equiv \lambda^* PH$, its consumption is given by

$$c_0^N = D^* - (R_0 D_0 - y_0).$$

Using the ((21)) and (C.3) we obtain

$$(1-\delta)(1-\gamma)C_0^N \cong (1-\delta)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}}\left(\left(c_0^N\right)^{\frac{\sigma-1}{\sigma}} + \beta\frac{\left(Y_0 - RD^*\right)^{\frac{\sigma-1}{\sigma}}}{\left(1-\delta^\sigma\hat{R}^{\sigma-1}\right)^{\frac{1}{\sigma}}}\right)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}} - 1.$$

Using (C.5) we can replace for human wealth Y_0 to obtain

$$\begin{aligned}
&(1-\delta)(1-\gamma)C_0^N \\
\cong\ &(1-\delta)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}}\left(\left(c_0^N\right)^{\frac{\sigma-1}{\sigma}} + \delta\frac{\left(D^* + \Delta D - \left[\begin{array}{l}\delta^\sigma\hat{R}^{\sigma-1}(R_0D_0 - y_0) \\ +\left(1-\delta^\sigma\hat{R}^{\sigma-1}\right)RD^*\end{array}\right]\right)^{\frac{\sigma-1}{\sigma}}}{1-\delta^\sigma\hat{R}^{\sigma-1}}\right)^{\frac{1-\gamma}{1-\frac{1}{\sigma}}} - 1.
\end{aligned}$$

The indifference equation (20) can then be written by equating C_0^N and C_0^I from this analysis. Note that this new indifference equation no longer contains the coefficient of relative risk aversion γ

except through the certainty-equivalent expected interest rate $\hat{R}$. In fact, this equation is *identical* to our indifference equation under the CRRA functional form, with

$$\tilde{R} = \frac{1}{1 - \beta^{\sigma} \hat{R}^{\sigma-1}},$$

i.e. different only in that future interest rates are replaced by their certainty equivalent.