

# Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches

MICHAEL CARLOS BEST

*Columbia University and NBER*

JAMES S. CLOYNE

*UC Davis, CEPR and NBER*

ETHAN ILZETZKI

*London School of Economics*

and

HENRIK J. KLEVEN

*Princeton University and NBER*

*First version received November 2017; Editorial decision October 2018; Accepted May 2019 (Eds.)*

Using a novel source of quasi-experimental variation in interest rates, we develop a new approach to estimating the Elasticity of Intertemporal Substitution (EIS). In the U.K., the mortgage interest rate features discrete jumps—notches—at thresholds for the loan-to-value (LTV) ratio. These notches generate large bunching below the critical LTV thresholds and missing mass above them. We develop a dynamic model that links these empirical moments to the underlying structural EIS. The average EIS is small, around 0.1, and quite homogeneous in the population. This finding is robust to structural assumptions and can allow for uncertainty, a wide range of risk preferences, portfolio reallocation, liquidity constraints, present bias, and optimization frictions. Our findings have implications for the numerous calibration studies that rely on larger values of the EIS.

*Key words:* Elasticity of Intertemporal Substitution, mortgage borrowing, interest rate notches, bunching, structural estimation

*JEL Codes:* D1, D14, D15, E2, E21, E4, E43, H3, H31

## 1. INTRODUCTION

How responsive are households to changes in the intertemporal price of consumption? In standard economic models, this response is governed by the Elasticity of Intertemporal Substitution (EIS). The EIS is a key parameter in economics as it plays a central role for a range of questions in macro, public finance, household finance, and asset pricing. Unfortunately, there exists no consensus on a reasonable value for this parameter due to limitations in data and research designs. The most

---

*The editor in charge of this paper was Adam Szeidl.*




FIGURE 1

Observed LTV distribution among U.K. refinancers

*Notes:* This figure shows the observed distribution of LTV ratios among refinancers in the U.K. between 2008–14. There are interest rate notches at LTV ratios of 60%, 70%, 75%, 80%, 85%, and 90% (depicted by vertical lines).

cited estimates in the literature range between 0 and 2, which is an enormous range in terms of its implications for intertemporal behaviour and policy.

A fundamental difficulty in addressing this question is how to find exogenous variation in interest rates. Most studies rely on time series movements in interest rates, which are gradual and almost certainly endogenous to unobserved factors that affect consumption. Our starting point is a novel source of quasi-experimental variation in interest rates arising from the fact that U.K. banks offer *notched* mortgage interest schedules. That is, the mortgage interest rate features discrete jumps at critical thresholds for the loan-to-value (LTV) ratio. For example, the interest rate increases by almost 0.5pp on the entire loan when crossing the 80% LTV threshold. This creates very strong incentives to reduce borrowing to a level below the notch, thereby giving up consumption today in order to get a lower interest rate and more consumption in the future. Intuitively, the magnitude of such borrowing and consumption responses to interest rate notches is governed by the value of the EIS.

Our study is based on administrative mortgage data from the Financial Conduct Authority (FCA). The data cover the universe of household mortgages in the U.K. between 2008 and 2014, including rich information on mortgage contracts and borrower characteristics. The majority of U.K. 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 strong incentives to refinance at the time the reset rate starts to apply. This makes refinancing a common occurrence in the U.K. We focus on the population of refinancers, because they allow for a clean assessment of borrowing and intertemporal consumption choices. Specifically, because housing choices are pre-determined for refinancers, estimating LTV responses in this sample allows us to isolate borrowing choices from housing choices.

Figure 1 plots the LTV distribution for U.K. home refinancers around the different interest rate notches, depicted by vertical lines. There is large and sharp bunching below every notch along with missing mass above every notch, which provides direct evidence that borrowers respond to interest rates. A recent literature in public economics has developed approaches to translate such bunching moments into reduced-form price elasticities, mostly focusing on behavioural responses

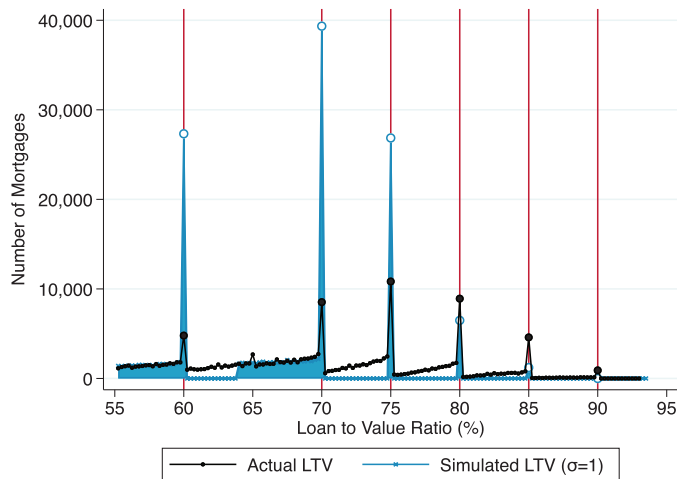


FIGURE 2

Observed versus simulated LTV distributions under an EIS = 1

*Notes:* This figure compares the observed LTV distribution (black markers) to a simulated LTV distribution (lighter markers) under Cobb–Douglas preferences (EIS = 1). The simulation is based on the standard lifecycle model introduced in Section 3 in which households choose their LTV optimally. Cobb–Douglas preferences are very far from being able to fit the data.

to taxes and transfers in static contexts (Saez, 2010; Chetty *et al.*, 2011; Kleven and Waseem, 2013). It remains an open question whether these bunching-based elasticities have any structural or external validity, and whether their interpretation is robust to allowing for dynamics (Kleven, 2016). In this article, we consider an inherently dynamic decision context and take the bunching literature in a more structural direction.

Translating bunching moments—or indeed any quasi-experimental moment—into structural parameters that can be used for out-of-sample prediction requires a structural model (see Einav *et al.*, 2015, 2017; Kleven, 2016). We develop two different approaches. The first approach is based on a simple two-period model with no uncertainty, no portfolio choice, no liquidity demand, and several other simplifying assumptions. This model provides the most transparent way of translating a bunching moment into the EIS. The second approach is based on a rich stochastic lifecycle model that relaxes many of the simplifying assumptions made in the baseline model. This model is more realistic, but computationally more involved and thus more of a “black box”. We show that these two approaches give essentially the same answer: the observed bunching at interest notches is consistent with a small EIS, around 0.1. We present a battery of robustness checks to and extensions of the stochastic lifecycle model, which confirm our finding that the EIS is small.

How can the EIS be small given the observation of substantial bunching? What would the raw data have looked like if the EIS were larger? To answer these questions, Figure 2 compares the observed LTV distribution to a simulated LTV distribution based on our two-period model with an EIS set equal to one (log preferences). These distributions are starkly different. The simulated distribution has zero mass above the 70% LTV threshold, except at the notches. The simulated distribution has *smaller* bunching at the highest notches (80% and 85%) and much *larger* bunching at the lowest notches (60%, 70%, and 75%). The fact that there is smaller bunching at the top reflects that, under an EIS of one, some homeowners jump across multiple notches and therefore skip the top notches entirely. We will show that the observed data are not

consistent with such multiple-notch jumps. The stark contrast between the observed and simulated distributions suggests that the data are inconsistent with standard assumptions about the EIS.

A key question for identification is whether we can close the gap between the observed and simulated distributions through other means than a small EIS. As we show in the article, it is not possible to close the gap by changing parameters and assumptions within a framework of frictionless household optimization. Rather, the only threat to identification is the presence of optimization frictions that attenuate bunching. Some borrowers may be stuck at LTV ratios above a notch, not because of true intertemporal preferences, but because they do not pay attention to or understand the incentives created by the notch. However, an important advantage of notch-based identification is that the missing mass just above the threshold is directly informative of optimization frictions. As shown by [Kleven and Waseem \(2013\)](#), it is possible to correct for frictions using missing mass in dominated regions just above notches. We develop a structural extension of the Kleven–Waseem friction approach, showing that optimization frictions in this setting are not sufficient to justify large values of the EIS.

Our article contributes to three literatures. First, we contribute to a large structural literature studying intertemporal substitution in consumption, reviewed by [Attanasio and Weber \(2010\)](#). This literature estimates consumption Euler equations using either aggregate data (*e.g.* [Hall, 1988](#); [Campbell and Mankiw, 1989](#)) or micro survey data (*e.g.* [Zeldes, 1989](#); [Attanasio and Weber, 1993, 1995](#); [Vissing-Jørgensen, 2002](#); [Gruber, 2013](#)). Most of the literature has relied on time series movements in interest rates, producing a very wide range of estimates depending on the analysis sample and empirical specification.<sup>1</sup> 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 rate changes over time, and that our estimating equation is not a standard Euler equation due to the discontinuous nature of the notched incentive. Our EIS estimates are at the lower end of the spectrum provided by these non-experimental studies.<sup>2</sup>

Second, we contribute to a reduced-form literature studying borrowing responses to the after-tax interest rate. This literature includes a number of natural experiment studies using variation in the after-tax interest rate created by taxes, subsidies, and regulation (*e.g.* [Follain and Dunskey, 1997](#); [Ling and McGill, 1998](#); [Dunskey and Follain, 2000](#); [Martins and Villanueva, 2006](#); [Jappelli and Pistaferri, 2007](#); [DeFusco and Paciorek, 2017](#)).<sup>3</sup> The range of estimates is very wide, from a zero effect in [Jappelli and Pistaferri \(2007\)](#) to elasticities of about 1 in [Dunskey and Follain \(2000\)](#) and 1.5–3.5 in [Follain and Dunskey \(1997\)](#). We estimate reduced-form borrowing elasticities around 0.5. A conceptual contribution of our article is to characterize the relationship between reduced-form borrowing elasticities and the structural EIS, demonstrating that the former by itself is not very informative about the latter. The translation between the two parameters is mediated by additional (endogenous) variables that can vary widely across borrower populations.<sup>4</sup>

Third, we contribute to the recent bunching literature in public economics (as reviewed by [Kleven, 2016](#)). Most of this literature has focused on static contexts and reduced-form estimation.

1. A methodological exception is [Gruber \(2013\)](#) who uses cross-sectional and time series variation in capital income tax rates to identify the EIS and obtains very large estimates of about 2.

2. [Havránek \(2015\)](#) conducts a meta-analysis of the existing literature and finds estimates centred around 0.3–0.4, after controlling for publication bias.

3. Related to our empirical approach, [DeFusco and Paciorek \(2017\)](#) estimate leverage responses using an interest notch created by the conforming loan limit in the U.S., although their estimates do not separate mortgage demand from housing demand as we do here. Most importantly, they do not pursue the analysis of structural parameters, which is the main contribution of our article.

4. This finding echoes insights from early calibration studies, which showed 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](#)).

By combining a bunching approach with dynamic structural estimation, our article is related to recent work by Einav *et al.* (2015, 2017) who analyse bunching at a kink point in U.S. Medicare. They argue that the choice of model is crucial when translating bunching into a parameter that can be used for out-of-sample prediction. In particular, they highlight the role played by frictions in the form of lumpiness and randomness in the choice variable used to bunch.<sup>5</sup> This contrasts with our finding that the structural EIS (“out-of-sample prediction”) is robust to the modelling assumptions we make. This difference can be explained mainly by a conceptual difference between kink-based and notch-based estimation. In the case of notches, the amount of friction is directly accounted using an observational moment—the amount of missing mass above the notch—as opposed making parametric assumptions about such frictions.

The article is organized as follows. Section 2 describes the institutional setting and data, Section 3 characterizes the link between bunching, reduced-form elasticities, and the EIS in our baseline two-period model, Section 4 presents empirical results using the baseline model, Section 5 develops and structurally estimates our full stochastic lifecycle model, and finally Section 6 concludes.

## 2. INSTITUTIONAL SETTING, DATA AND DESCRIPTIVES

### 2.1. U.K. mortgage market

The U.K. mortgage market has several institutional features that facilitate our analysis. First, the interest rate on mortgage debt follows a step function with discrete jumps—notches—at certain LTV thresholds. There are interest rate notches at LTVs of 60%, 70%, 75%, 80%, and 85%.<sup>6</sup> When a borrower crosses one of these thresholds, the interest rate increases on the entire loan. The thresholds apply to the LTV ratio at the time of loan origination; the interest rate does not change as amortization or house price growth gradually reduces the LTV. The size of the interest rate jump at a given threshold varies across product types and over time.<sup>7</sup> The notches are very salient: daily newspapers display menus of interest rates by bank and LTV bracket, and the LTV thresholds feature very prominently when shopping for mortgages. For example, the mortgage websites of all the major banks show LTV brackets and interest rates for their different products up front.<sup>8</sup>

Second, most U.K. mortgage products come with a relatively low interest rate for an initial period—typically 2, 3, or 5 years—after which a much larger (and variable) reset rate starts to

5. As discussed by Kleven (2016), this general insight echoes findings elsewhere in the bunching literature showing that the conversion of observed bunching (at kinks) into a structural elasticity is very sensitive to the assumed model of optimization frictions (*e.g.* Saez, 1999; Chetty *et al.*, 2010, 2011; Gelber *et al.*, 2019).

6. There is in principle also an interest notch at 90%. However, very few banks offered mortgages above 90% after the financial crisis, implying that this threshold became a corner solution rather than a notch for most borrowers in our data. Our empirical analysis therefore focuses on the notches below 90%.

7. There is also some—but much less—variation in the size of notches across banks within product type and time. In particular, some banks do not feature certain notches at some points in time, but we show later that such no-notch observations represent a very small fraction of the data.

8. A broad question not addressed in this article is *why* U.K. banks impose such notched interest rate schedules, a type of question that often arises in settings with notched incentive schemes (Kleven, 2016). The traditional explanation for upward-sloping interest rate schedules is that the default risk is increasing in leverage, either due to increasing risk for each borrower or due to adverse changes in the mix of borrowers. However, under the reasonable assumption of *smoothly* increasing default rates, standard models predict smoothly increasing interest rates. While the U.K. practice of implementing the increasing interest rate schedule as a step function may not be second-best efficient in standard models, it may be explained—as with other types of notches—by the simplicity and salience of notches to banks and their customers. Our empirical analysis of these notches is implicitly based on the assumption that default rates (in the absence of notches) are smooth around the threshold.

apply. The notched interest rate schedule described above applies to the rate charged during the initial period of 2–5 years as opposed to the rate charged over the entire term of the mortgage (typically 25–35 years). The large and variable reset rate creates a very strong incentive to refinance at the end of the initial lower-rate period. This makes refinancing a frequent occurrence in the U.K. In this article, we focus specifically on refinancers as this will allow us to isolate borrowing choices from housing choices, which is critical when assessing intertemporal consumption substitution.

Third, while borrowers have a strong incentive to refinance no later than at the onset of the reset rate, the cost of early refinancing means that there is also a strong incentive to refinance no sooner than this time. Specifically, U.K. mortgage contracts feature large prepayment charges (often 5–10% of the outstanding loan) on borrowers who refinance before reset rate onset. The combination of penalizing reset rates and heavy prepayment charges implies that households have strong incentives to refinance right around the end of the initial lower interest period.

To confirm that households act on these refinancing incentives, [Supplementary Appendix Figure A.1](#) shows the distribution of time-to-refinance in our data. The distribution features large excess mass in refinancing activity around 2, 3, and 5 years, consistent with the fact that these are the most common timings of the penalizing reset rate. The lightly shaded bars indicate the fraction of households in each monthly bin who refinance “on time”, *i.e.* around the time their reset rate kicks in. These bars show that the majority of households refinance around reset rate onset and that this can explain the excess mass at 2, 3, and 5 years.<sup>9</sup> Note that this graph represents clear evidence that borrowers respond to interest rate changes, but on a different margin (refinance timing) than our main focus (borrowing and consumption). What is more, the empirical patterns documented here imply that the time of refinancing is effectively locked in by the reset rate structure. This is helpful for ruling out selection issues from endogenous refinance timing in the analysis below.

## 2.2. Data

Our analysis uses a novel and comprehensive administrative dataset containing the universe of mortgage product sales in the U.K.<sup>10</sup> This Product Sales Database (PSD) is collected by the FCA for regulatory purposes and has information on mortgage originations back to April 2005. This includes detailed information on the mortgage contract 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 higher reset rate starts applying), whether the mortgage payments include amortization, and the mortgage term over which the full loan will be repaid. The data also include a number of borrower characteristics such as age, income, whether the income is solely or jointly earned, whether the borrower is a first-time buyer, mover or refinancer, and the reason for the refinance. There are also some characteristics of the property such as the type of dwelling and the number of rooms.

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. For example, it is possible to observe an actual LTV ratio of 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 U.K. available from the organization

9. Indeed, we observe only around a quarter of households not refinancing on time.

10. The FCA PSD covers regulated mortgage contracts, which represent the majority of home finance products in the U.K. The database does not cover home finance products such as home purchase plans, home reversions, second-charge lending, and buy-to-let mortgages.

MoneyFacts between 2008Q4 and 2014Q4. For a mortgage observation in the PSD, we find the corresponding product in MoneyFacts based on the lender, the date of the loan, the mortgage type, and the interest rate. Where the interest rate paid accords with an LTV bracket just below the actual LTV in the PSD data, we subtract the product fee from the observed loan value. Inspecting such individuals, the loan amount in excess of the threshold often corresponds precisely to the product fee. As a result, this adjustment places most of these 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 refinancing. Using information about the characteristics of the property and the borrower, we can match refinancers over time in order to construct a panel.<sup>11</sup> As described later, the panel structure allows us to implement a novel approach for estimating the counterfactual LTV distribution absent notches. The refinancer panel will therefore be the baseline dataset for our analysis.

Table 1 shows a 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 contains around 2.8 million observations. Column 2 shows how the sample characteristics change when we restrict attention to refinancers. The descriptive statistics are very similar, although the LTV and loan-to-income (LTI) ratios are slightly lower for refinancers as one would expect. Column 3 shows the descriptive statistics in the panel of refinancers that we use in the empirical analysis. In moving from column 2 to 3, we lose refinancers for whom we lack sufficient information on 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 to the full population of borrowers.

### 2.3. *Interest rate jumps at notches*

As described above, the U.K. 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 estimate the size of these interest rate notches. Unlike standard bunching approaches in which the discontinuity is the same across agents, in our setting the interest rate notch varies by bank, mortgage product, and the time of loan origination (all of which we observe). As we will show, notches do not depend on individual characteristics conditional on bank, product, and time, which is important for ruling out selection bias in the estimated interest notches. This is because the U.K. mortgage market works like a mortgage supermarket in which banks offer their interest rate schedule on a given product to all borrowers who meet their lending standards, as opposed to entering into individual negotiations that depend on idiosyncratic factors.

Our empirical analysis will be based on the average interest rate jump at each notch conditional on bank, product, and time. We estimate these interest rate jumps non-parametrically using the following regression:

$$r_i = f(\text{LTV}_i) + \beta_1 \text{bank}_i + \beta_2 \text{variability}_i \otimes \text{duration}_i \otimes \text{month}_i + \beta_3 \text{repayment}_i + \beta_4 \text{term}_i + v_i, \quad (2.1)$$

where  $r_i$  is the nominal mortgage interest rate for individual  $i$ ,  $f(\cdot)$  is a step function with steps at each 0.25pp of the LTV ratio,  $\text{bank}_i$  is a vector of bank dummies,  $\text{variability}_i$  is a vector of interest

11. For each homeowner we use the location of their house by 6-digit postcode (a code that covers a very small geographical area, around 15 homes on average) and the date of birth of the homeowner.



TABLE 1  
Descriptive statistics across samples

	(1) Full sample mean/sd	(2) Refinancer mean/sd	(3) Refinancer panel mean/sd
Interest rate (%)	4.12 (1.29)	4.26 (1.38)	4.06 (1.23)
Loan size (£)	142,710.6 (120,512.8)	138,376.2 (120,725.2)	145,400.8 (116,489.4)
Property value (£)	257,004.0 (255,284.2)	263,076.4 (270,122.1)	262,304.0 (247,074.1)
Loan-to-value ratio (%)	59.9 (22.1)	56.4 (21.7)	58.8 (19.8)
Gross income (£)	57,270.6 (84,069.3)	57,480.9 (84,517.4)	57,483.2 (77,513.0)
Loan-to-income ratios	2.78 (1.85)	2.68 (2.03)	2.78 (1.60)
Repayments to income (%)	21.8 (39.5)	22.6 (48.4)	21.9 (15.4)
Proportion with joint income	0.54 (0.50)	0.55 (0.50)	0.55 (0.50)
Mortgage term (years)	20.7 (7.44)	19.0 (6.92)	20.0 (6.65)
Time to refinance (years)	2.80 (1.71)	2.80 (1.70)	3.27 (1.72)
Proportion of fixed rate mortgages	0.68 (0.47)	0.65 (0.48)	0.73 (0.45)
Proportion of refinance events	0.56 (0.50)	0.88 (0.32)	0.76 (0.42)
Borrower's age	40.0 (10.0)	41.7 (9.57)	40.3 (8.93)
Observations	3,049,164	1,961,325	647,192

*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 refs) including (a) observations where we can find information on the product fee in MoneyFacts (b) any refinancer's previous mortgage (which we can use without fee information). Column (2) restricts the sample to refinancers only (all refinance products and any refinancer's previous mortgage). Finally, column (3) shows the sample of refinancers in the panel (refinancers we can link over time and where we can construct a counterfactual LTV ratio). The samples in columns (2) and (3) use all mortgage events for refinancing households. This includes each household's first mortgage which, by definition, was not a refinance. As a result, the proportion of refinance events in columns (2) and (3) is not equal to 1.

variability dummies (fixed interest rate, variable interest rate, capped interest rate, and "other"),  $\text{duration}_i$  is a vector of dummies for the duration of the initial low interest period (the time until the reset rate kicks in),  $\text{month}_i$  is a vector of dummies for the month in which the mortgage was originated,  $\text{repayment}_i$  is a vector of dummies for the repayment type (interest only, capital and interest, and "other"), and  $\text{term}_i$  is a vector of dummies for the total term length. We denote by  $\otimes$  the outer product, so that the term  $\text{variability}_i \otimes \text{duration}_i \otimes \text{month}_i$  allows each combination of interest rate variability and duration to have its own non-parametric time trend.

Figure 3 plots the conditional interest rate as a function of LTV based on specification (2.1). 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%. These interest jumps are larger at LTV thresholds higher up in the distribution. At the two top thresholds, the annual interest rate increases by almost 0.5pp. Importantly, the interest rate schedule is flat between notches. This






FIGURE 3

## Interest rate jumps at notches

*Notes:* This figure shows the conditional interest rate as a function of the LTV ratio from the non-parametric regression (2.1). In each LTV bin, we plot the coefficient on the LTV bin dummy plus a constant given by the mean predicted value  $E[\hat{r}_i]$  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 notches at LTV ratios of 60%, 70%, 75%, 80%, and 85%.

implies that, conditional on product and bank characteristics, the mortgage interest rate is almost fully determined by the LTV notches we exploit.

The flatness of the interest schedule between notches suggests that individual characteristics (that vary by LTV) have no effect on the mortgage interest rate. [Supplementary Appendix Figure A.2](#) verifies this by controlling for the individual characteristics we observe (such as age, income, and family status) in the estimation of the interest schedule. The figure shows that the results are virtually unchanged. If observables such as age and income do not matter for the interest notches, it is difficult to imagine any unobservables that would matter. These results confirm the institutional context described earlier, namely that the U.K. mortgage market works as a mortgage supermarket in which a given type of product is offered at a given price, independently of who buys it.<sup>12</sup>

When estimating the interest jumps from the coefficients on the LTV bin dummies in equation (2.1), we are holding all non-LTV mortgage characteristics constant on each side of the LTV threshold. For example, if a household is observed in a 5-year fixed rate 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 a 5-year fixed to a 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

12. Moreover, the *global* interest estimations shown in Figure 3 and [Supplementary Appendix Figure A.2](#) understate flatness between notches compared to the more precise local estimations used later. The locally estimated interest schedules are essentially completely flat. This implies that “donut hole” approaches in which we exclude observations in a range around the threshold when estimating the interest rate jump give virtually unchanged results.

LTV bin, if lower interest rate products or banks are available, in equilibrium this must be offset by less favourable terms in other dimensions. In this case, the within-product interest rate jump around the threshold is the right measure of the price incentive.

#### 2.4. *Actual versus counterfactual LTV distributions*

The interest rate notches described above create strong incentives for borrowers to choose LTVs just below one of the notches, giving rise to bunching below the critical LTV thresholds and missing mass (holes) above them. We have already seen in Figure 1 that bunching and missing mass are indeed features of the data. The idea of our approach is to use these empirical moments to identify the EIS.

To quantify 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 (see Kleven, 2016). The standard approach is to fit a flexible polynomial to the observed distribution, excluding data around the notch, and then extrapolate the fitted distribution to the notch (Chetty et al., 2011; Kleven and Waseem, 2013). However, this approach is not well-suited for our context: it is 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 1 suggests that most 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 refinancer data. Based on the LTV in the previous mortgage, the amortization schedule, and the house value at the time of refinance (which is assessed by the bank), we measure the new LTV *before the refinancer has taken any action*. We label this the *passive LTV* as it would be the LTV if the homeowner simply rolled over debt between the two mortgage contracts. We will base our estimate of the counterfactual LTV on the passive LTV with an adjustment that we describe below.

In Figure 4A, we compare the actual LTV distribution to the passive LTV distribution. We see that the passive LTV distribution is smooth: unlike the actual LTV distribution it features no excess bunching or missing mass around notches. In general, the two distributions in Figure 4A may be different for two reasons: (i) behavioural 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.<sup>13</sup> Supplementary Appendix 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 LTV using the profile of equity extracted in Supplementary Appendix Figure A.3. The assumption we are making is that the equity extraction profile among non-bunchers is a good proxy for the equity extraction profile in the full population of refinancers (including bunchers) in the counterfactual scenario without notches. We relax this assumption below.

13. Where we define bunchers as those with LTVs within 0.25pp below a notch, and non-bunchers as their complement.

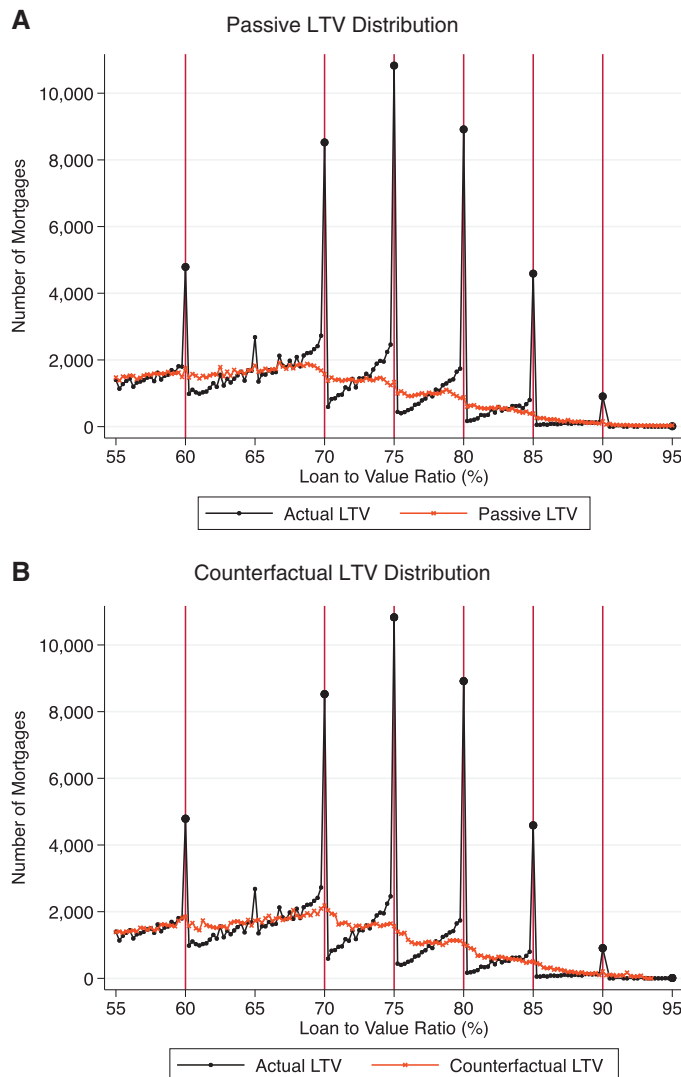


FIGURE 4

Constructing the counterfactual LTV distribution.

*Notes:* This figure shows the two steps in the construction of the counterfactual LTV distribution among refinancers. Each panel shows the actual LTV distribution with dots (as in Figure 1). Panel A shows the distribution of passive LTVs with crosses, calculated based on the LTV of the previous mortgage, amortization, and the house value at the time of refinancing. Panel B shows the distribution of counterfactual LTVs with crosses, which adjusts passive LTVs for the average equity extraction of non-bunchers in the actual distribution.

Our estimate of the counterfactual LTV distribution is shown in Figure 4B. Comparing the actual and counterfactual LTV distributions 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.

The assumption that the equity extraction profile among non-bunchers is a good proxy for the counterfactual equity extraction profile among bunchers raises potential concerns about selection. It is possible that bunchers are selected on variables that would impact their counterfactual equity extraction. We can address this concern in two ways. First, a straightforward extension of our approach is to control for selection on observables: income, age, family status, and the number of past and future bunching episodes. The last of these covariates intends to capture the possibility that bunchers at time  $t$  may be a selected sample of “optimizers” (thus bunching more at other times as well) while non-bunchers may be a selected sample of “passives”.<sup>14</sup> We regress equity extracted among non-bunchers on these covariates and predict equity extraction for both bunchers and non-bunchers from this regression. This approach makes virtually no difference to any of our results. Hence, if selection were an issue for our equity extraction adjustment, it would have to come from unobservables that impact equity extraction *and* are uncorrelated with (and therefore not picked up by) the observables that we do control for. Second, to allow for selection on unobservables, we can use a standard Heckman (1979) sample selection framework to estimate equity extraction. A previous version of the paper (Best *et al.*, 2018) considered such an extension and it had very little impact on our results.

It is worth pointing out that there is a very simple reason why the counterfactual distribution is robust to different ways of doing the equity extraction adjustment. The reason is that the adjustment corresponds to shifting a distribution—the passive LTV distribution—which is relatively flat, or at least not strongly sloped, around the notches. Of course, if the passive LTV distribution had been completely flat, any shift to the left or right would have precisely zero impact on the bunching estimation. This is not true here, but the passive LTV distribution is sufficiently flat that the specific procedure we use is relatively unimportant.

Finally, we note that the observed LTV distribution features a small spike at an LTV of 65%, although this threshold is not associated with an interest notch. This spike is most naturally explained by round-number bunching, a phenomenon observed across a wide range of settings (see Kleven, 2016). If we do not adjust for round-number bunching, the amount of excess mass at interest notches (all of which are located at round numbers) would overstate the true response to interest rates. While we could adjust for round-number bunching using the observed spike at 65%, a concern may be that round-number bunching is different in different parts of the LTV distribution. Instead, we deal with this issue by exploiting the fact that some banks at some points in time do not feature a specific notch. This allows us to net out round-number bunching at a given notch using bunching at that same threshold in a no-notch subsample. As we show in Section 4, this adjustment has only a minor impact on our results.

### 3. A SIMPLE STRUCTURAL MODEL

In this section, we develop an approach to estimating the EIS using bunching at interest rate notches. The approach is based on a two-period model in which we make many simplifying assumptions. The virtue of this model is to provide a simple and transparent mapping between an observed bunching moment and the underlying structural EIS. In Section 5, we show that the estimates are robust to extending the analysis to a rich stochastic lifecycle model.

14. However, it turns out that the number of previous/future bunching events is not excessively large for households currently bunching. Supplementary Appendix Figure A.4 shows the average number of past/future bunching events at each value of current chosen LTV. The graph is smoothly increasing and features no spikes at notches. This suggests that bunching households are not different “types” in terms of their general propensity to bunch or optimize.

### 3.1. The mapping between bunching and the EIS

We consider households who live for two periods (0 and 1) and have perfect foresight. They are homeowners and have chosen to remain in their current dwellings in both periods, but face a mortgage refinancing choice at time zero. As a baseline, assume that they can refinance at a constant gross borrowing rate equal to  $R$  (*i.e.* there is no notch).

The utility of consuming housing services  $H_t$  is separable from the utility of consuming non-durable goods  $c_t$ , and households place no value on residual wealth (*e.g.* bequests) at the end of period 1. Households value non-housing consumption in any period  $t$  via a constant EIS function  $\frac{\sigma}{\sigma-1} c_t^{\frac{\sigma-1}{\sigma}}$  and discount the future by a factor  $\delta$ . Hence, the lifetime utility derived from non-housing consumption is given by  $\frac{\sigma}{\sigma-1} \left( c_0^{\frac{\sigma-1}{\sigma}} + \delta c_1^{\frac{\sigma-1}{\sigma}} \right)$ .

The households receive an exogenous stream of income,  $y_t$  in period  $t$ . They have initial net wealth  $W_0$  equal to housing wealth net of any mortgage debt and net of any refinancing costs incurred in period zero. For simplicity, we assume that households hold no assets other than housing and have no liabilities other than the mortgage. The budget constraint in period 0 is therefore given by

$$c_0 = y_0 + W_0 - (1 - \lambda)P_0H, \quad (3.2)$$

where  $\lambda$  is the LTV of the new mortgage and  $P_0H$  is housing value (using that  $H_0 = H_1 = H$ ). The period-1 budget constraint is given by

$$c_1 = y_1 - R\lambda P_0H + (1 - d)P_1H, \quad (3.3)$$

where  $d$  is the rate of house depreciation and  $P_1$  is the house price in period 1.

Households choose consumption according to the standard Euler equation

$$c_1 = (\delta R)^\sigma c_0. \quad (3.4)$$

Equations (3.2)–(3.4) determine the choice of  $c_0$ ,  $c_1$ , and  $\lambda$  as functions the exogenous parameters of the model. We note that the LTV choice  $\lambda$  is monotonically decreasing in initial wealth  $W_0$  and the interest rate  $R$ .<sup>15</sup>

To begin with, we simplify by assuming that households are heterogeneous only in  $W_0$ . Our general argument goes through if households are heterogeneous in other dimensions such as income, housing quality, or preferences. Below we analyse the important case where the EIS parameter itself is heterogeneous. If  $W_0$  is smoothly distributed in the population, equations (3.2)–(3.4) imply a smooth density distribution of LTV, which we denote by  $f_0(\lambda)$ . We will refer to this as the *counterfactual* LTV distribution under a constant interest rate  $R$ . Our estimate of the empirical counterpart to  $f_0(\lambda)$  was shown in Figure 4.

Suppose now that an interest rate notch is introduced at  $\lambda^*$ , so that the borrowing rate jumps from  $R$  to  $R + \Delta R$  for LTVs exceeding  $\lambda^*$ . Figure 5 illustrates the implications of this notch for borrowing and consumption. Panel A depicts the period-1 budget constraint before and after the

15. The first effect follows from the fact that consumption in period 1 is a normal good and therefore increasing in initial wealth  $W_0$ . Initial wealth can increase  $c_1$  only via a decrease in borrowing in period 0. The second effect follows from the fact that the wealth and substitution effects of the interest rate push in the same direction here. Given that the household is a borrower, an increase in the interest rate reduces lifetime wealth and thus consumption in both periods. Reducing consumption in period 0 requires a reduction in debt. The substitution effect follows from the fact that an increase in the interest rate increases the relative price of consumption in period 0 (or equivalently the relative cost of debt).

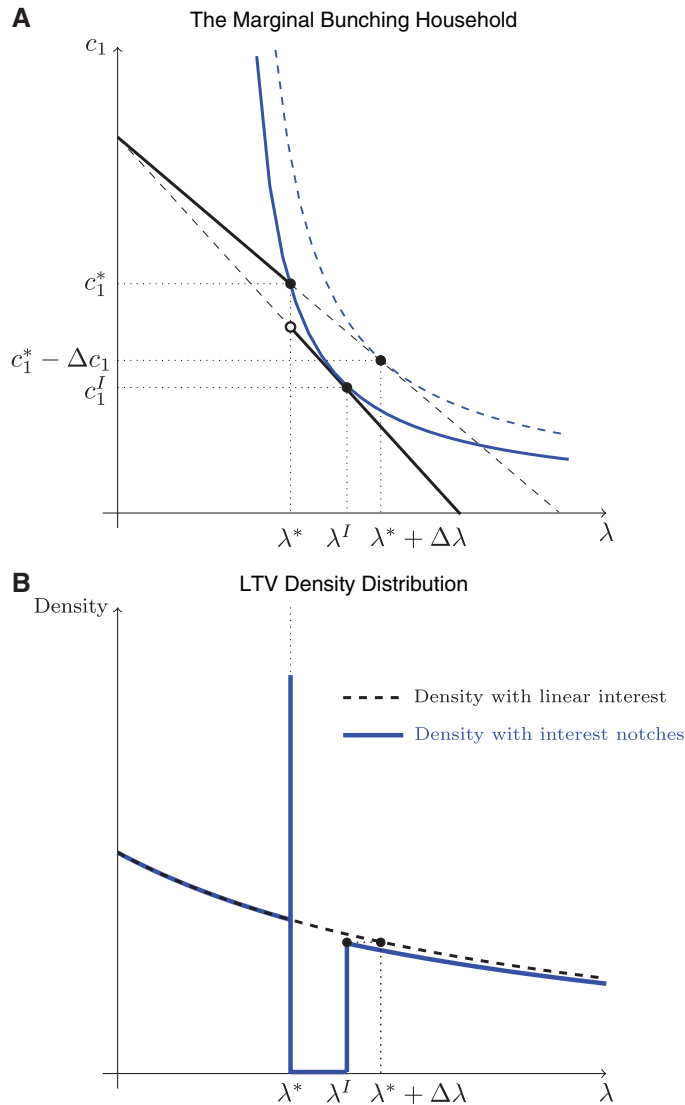


FIGURE 5

Bunching in a simple model of intertemporal substitution.

*Notes:* The figure shows the choice faced by a refinancing household faced with a notched interest rate schedule in a budget set diagram with LTV  $\lambda$ , on the horizontal axis and consumption in the future,  $c_1$ , on the vertical axis. Given the period zero budget constraint and the value of housing, an LTV choice  $\lambda$  translates into current consumption  $c_0$ . The solid line depicts the budget constraint with an interest notch at the LTV  $\lambda^*$ . The budget constraint has a slope  $-R$  below  $\lambda^*$  and a slope  $-(R + \Delta R)$  above  $\lambda^*$ . There is a discrete jump in the budget set (rather than a kink) because the interest on the entire loan jumps discretely at the notch. The indifference curves shown are those of the marginal bunching household that is indifferent between bunching at the notch and borrowing  $\lambda^*$ , and a point  $\lambda^I$  in the interior of the higher interest rate bracket. This household would have chosen a point  $\lambda^* + \Delta\lambda$  in the absence of the interest rate notch. As described in Section 3, the optimality conditions for  $\lambda^I$  and  $\lambda^* + \Delta\lambda$  together with the household's indifference between  $\lambda^I$  and  $\lambda^*$  identify the curvature of the indifference curves—the EIS—in terms of observable and/or estimable quantities.

introduction of the notch in  $\{\lambda, c_1\}$  space. It also shows the indifference curves before and after the notch for the *marginal bunching household*, i.e. the highest-LTV (lowest-wealth) household who will choose to bunch at the notch.<sup>16</sup> When faced with the constant interest rate  $R$ , this household chooses an LTV of  $\lambda^* + \Delta\lambda$ , where the indifference curve is tangent to the initial budget constraint. After the introduction of the notch, this household is indifferent between locating at the LTV threshold  $\lambda^*$  and locating at the best interior LTV  $\lambda^I$ , where the indifference curve is tangent to the notched budget constraint. All households whose LTV fell in the segment  $[\lambda^*, \lambda^* + \Delta\lambda]$  absent the notch are strictly better off bunching than staying at an interior LTV.

Panel B shows the LTV distribution before and after the notch. In the presence of the notch, there is sharp bunching at  $\lambda^*$  along with a hole in the distribution between  $(\lambda^*, \lambda^I)$ . The amount of bunching is equal to  $B = \int_{\lambda^*}^{\lambda^* + \Delta\lambda} f_0(\lambda) d\lambda \simeq f_0(\lambda^*) \Delta\lambda$ . Hence, with estimates of excess bunching  $B$  and the counterfactual density around the notch  $f_0(\lambda^*)$ , it is possible to estimate the LTV response  $\Delta\lambda$ . The fundamental idea of our approach—a dynamic extension of [Kleven and Waseem \(2013\)](#)—is that we can use the indifference condition between  $\lambda^*$  and  $\lambda^I$  for the marginal buncher to derive a relationship between the LTV response  $\Delta\lambda$  and the EIS  $\sigma$ .

To characterize the estimating indifference equation, we first use that the marginal bunching household chooses the LTV ratio  $\lambda^* + \Delta\lambda$  in the counterfactual scenario with a constant interest rate  $R$ . From equations (3.2)–(3.4), this allows us to relate initial wealth  $W_0$  for this household to the other parameters of the model as follows:

$$W_0 = P_0 H - y_0 + \frac{y_1 + (1-d)P_1 H - ((\delta R)^\sigma + R)(\lambda^* + \Delta\lambda)P_0 H}{(\delta R)^\sigma}. \quad (3.5)$$

This relationship allows us to eliminate  $W_0$  from the problem. This is helpful because our data do not contain information on non-housing assets and liabilities, and therefore do not enable us to measure total initial wealth.

Using wealth defined in equation (3.5) and the optimality conditions (3.2)–(3.4) evaluated at the interest rate  $R + \Delta R$ , we can solve for the lifetime utility of the marginal buncher at the best interior choice  $\lambda^I$  in the presence of the notch. This is given by

$$V^I(\sigma, \delta, \Delta\lambda, \Delta R, \mathbf{x}) = \frac{\sigma}{\sigma - 1} (P_0 H)^{\frac{\sigma-1}{\sigma}} \frac{(\delta^\sigma (R + \Delta R)^{\sigma-1} + 1)^{\frac{1}{\sigma}}}{(\delta R)^{\sigma-1}} \times \left( \left( \frac{(\delta R)^\sigma}{R + \Delta R} + 1 \right) \left( \frac{y_1}{P_0 H} + \Pi_1 \right) - ((\delta R)^\sigma + R)(\lambda^* + \Delta\lambda) \right)^{\frac{\sigma-1}{\sigma}}, \quad (3.6)$$

where  $\Pi_1 \equiv (1-d)\frac{P_1}{P_0}$  is gross house price growth net of depreciation. In the indirect utility function  $V^I(\cdot)$ , the argument  $\mathbf{x}$  is a vector that includes the parameters  $\{\lambda^*, R, \frac{y_1}{P_0 H} + \Pi_1\}$ .

Similarly, setting  $\lambda = \lambda^*$  and applying the interest rate  $R$ , the budget constraints (3.2) and (3.3) and the wealth condition (3.5) allow us to evaluate lifetime utility at the notch as

$$V^N(\sigma, \delta, \Delta\lambda, \mathbf{x}) = \frac{\sigma}{\sigma - 1} (P_0 H)^{\frac{\sigma-1}{\sigma}} \left( \frac{1}{(\delta R)^\sigma} \left( \frac{y_1}{P_0 H} + \Pi_1 - R\lambda^* - ((\delta R)^\sigma + R)\Delta\lambda \right)^{\frac{\sigma-1}{\sigma}} + \delta \left( \frac{y_1}{P_0 H} + \Pi_1 - R\lambda^* \right)^{\frac{\sigma-1}{\sigma}} \right). \quad (3.7)$$

The marginal buncher is indifferent between bunching at the notch and locating at the best interior LTV, allowing us to state the following proposition:

16. Indifference curves can be plotted in  $\{\lambda, c_1\}$  space using the period-0 budget constraint.



**Proposition 1 (Estimating indifference equation)** Given a bunching moment  $\{\Delta\lambda, \Delta R\}$  and a discount factor  $\delta$ , the EIS  $\sigma$  is the solution to the indifference equation

$$F(\sigma, \delta, \Delta\lambda, \Delta R, \mathbf{x}) \equiv V^N(\sigma, \delta, \Delta\lambda, \mathbf{x}) - V^I(\sigma, \delta, \Delta\lambda, \Delta R, \mathbf{x}) = 0, \quad (3.8)$$

where  $\mathbf{x} = \left\{ R, \lambda^*, \frac{y_1}{P_0 H} + \Pi_1 \right\}$ , and where  $V^I(\cdot)$  and  $V^N(\cdot)$  are given by (3.6) and (3.7), respectively.

*Proof.* The proof is in [Supplementary Appendix B](#).  $\parallel$

Three points are worth highlighting. First, the indifference equation (3.8) is based on a setting with only one notch, while our empirical setting has multiple notches. In the presence of multiple notches, it is possible that bunchers move across more than one threshold at a time, and it is conceptually straightforward to modify the indifference equation to allow for this (see [Kleven and Waseem, 2013](#)). We focus on the single-notch equation here, because the data does not support the presence of multiple-notch jumps. We discuss this point in the next section.

Second, the indifference equation contains two structural parameters, the EIS  $\sigma$  and the discount factor  $\delta$ . This suggests that we cannot identify  $\sigma$  from a single bunching moment. However, it turns out that the value of  $\sigma$  is extremely robust to assumptions about  $\delta$ . The intuitive reason is that the discount factor primarily governs the *level* of borrowing at any interest rate (*i.e.* it shifts the LTV distribution on both sides of the notch) and has only a minor impact on the *response* of borrowing to interest rate changes. In contrast, the EIS governs the curvature of intertemporal preferences, which directly impacts bunching responses. This is immediately apparent when differentiating the Euler equation (3.4):

$$\frac{\partial \log(c_1/c_0)}{\partial \log R} = \sigma. \quad (3.9)$$

The response of consumption growth to a small change in the interest rate is fully governed by the EIS with no role for the discount factor. Of course, the bunching moment reflects the response of borrowing rather than consumption, but these two are intimately related as shown by the budget constraints (3.2) and (3.3). Furthermore, the bunching moment reflects a response to a discontinuous, rather than marginal, interest change, which implies that the Euler equation logic does not carry over exactly. Still, a similar logic implies that the discount factor plays a very small role.<sup>17</sup> In Section 3.3, we demonstrate this important identification argument using numerical simulations.<sup>18</sup>

Third, even in this simple dynamic model, the estimating indifference equation is considerably more involved than the static bunching estimator developed by [Kleven and Waseem \(2013\)](#). The static bunching estimator does not require calibrating any variables: the bunching moment maps directly into a structural elasticity. The added complexity of the dynamic approach increases by an order of magnitude when we turn to the full stochastic lifecycle model in Section 5. However,

17. A formal demonstration of the relative importance of the EIS and the discount factor can be obtained by differentiating the indifference equation (3.8) with respect to  $\sigma$  and  $\delta$ . However, because of the complexity of the indifference equation, the resulting expressions are not helpful for intuition and yield insight only with additional numerical assumptions. An exercise of this sort shows that  $\sigma$ 's effect on bunching is orders of magnitude larger than  $\delta$ 's effect.

18. Besides  $\delta$ , the only other value in equation (3.8) that requires calibration is  $\frac{y_1}{P_0 H} + \Pi_1 = \frac{y_1 + (1-d)P_1 H}{P_0 H}$ . This is a measure of future resources from human wealth ( $y_1$ ) and housing wealth  $((1-d)P_1 H)$ , scaled by current housing wealth. We estimate  $\sigma$  using empirically reasonable values of this variable, but results are essentially unaffected by assumed parameter values (as for  $\delta$ , this is mainly a level effect rather than a response effect).

as we will show, it is a general feature of our methodology that the calibrated variables have a very small impact on the estimating indifference equation, making our results robust despite the analytical complexity of the expressions. The intuitive reason is essentially the same as the one underlying the robustness to  $\delta$ .

The exposition above assumes that there is only one value of the structural EIS  $\sigma$ , while in practice there is likely to be heterogeneity in this parameter. In fact, the empirical LTV distribution shown in Figure 1 implies that this has to be the case: without heterogeneity, there would be a *sharp* hole in the LTV distribution between  $\lambda^*$  and  $\lambda^I$  as illustrated in Figure 5B, whereas the empirical LTV distribution features a gradual hole and has some refinancers located just above the notch.<sup>19</sup> This provides *prima facie* evidence that some households have very small  $\sigma$ s while others have larger  $\sigma$ s. As Kleven and Waseem (2013) and Kleven (2016) show, in the presence of heterogeneity in  $\sigma$ , our bunching approach estimates the average  $\sigma$ .

To see this, consider a joint distribution of initial wealth  $W_0$  and the EIS  $\sigma$ . At each elasticity level  $\sigma$ , households optimize as characterized above. In the counterfactual scenario with a constant interest rate  $R$ , there is a joint distribution of LTV and EIS given by  $g_0(\lambda, \sigma)$  and an unconditional distribution of LTV given by  $g_0(\lambda) = \int_{\sigma} g_0(\lambda, \sigma) d\sigma$ . In the observed scenario with a notched interest rate, the marginal buncher at elasticity level  $\sigma$  reduces LTV by  $\Delta\lambda_{\sigma}$ . We can then link bunching  $B$  to the average LTV response at the notch  $E[\Delta\lambda_{\sigma}|\lambda^*]$  as follows

$$B = \int_{\sigma} \int_{\lambda^*}^{\lambda^* + \Delta\lambda_{\sigma}} g_0(\lambda, \sigma) d\lambda d\sigma \simeq g_0(\lambda^*) E[\Delta\lambda_{\sigma}|\lambda^*], \quad (3.10)$$

where the approximation assumes that the counterfactual density  $g_0(\lambda, \sigma)$  is roughly constant in  $\lambda$  on the bunching segment  $(\lambda^*, \lambda^* + \Delta\lambda_{\sigma})$ . In other words, in the presence of heterogeneous treatment effects, bunching identifies a local average treatment effect. When applying a bunching moment like  $E[\Delta\lambda_{\sigma}|\lambda^*]$  to the estimating indifference equation (3.8), we are estimating EIS *at the average LTV response* as opposed to the average EIS. These two will in general be different due to the non-linearity of (3.8), creating a form of aggregation bias. As elaborated by Kleven (2016), such aggregation bias is likely to be very small in practice.

A large literature estimates reduced-form elasticities of borrowing or saving with respect to the interest rate. How does one compare the magnitude of such reduced-form elasticities to the EIS? We can use our framework to characterize the relationship between the two elasticity concepts. Denoting the elasticity of borrowing with respect to the interest rate by  $\varepsilon$ , comparative statics on (3.2)–(3.4) give the following result.

**Proposition 2 (EIS versus reduced-form borrowing elasticity)** Given the EIS  $\sigma$ , the discount factor  $\delta$ , the gross interest rate  $R$ , and the ratio  $\text{LTW} \equiv \frac{P_0 H - W_0 - y_0}{y_1 + (1-d)P_1 H}$ , the elasticity of borrowing with respect to the interest rate is given by

$$\varepsilon = -\frac{\partial \log \lambda}{\partial \log R} = \frac{\sigma(\delta R)^{\sigma} + R}{(\delta R)^{\sigma} + R} - \frac{\sigma(\delta R)^{\sigma} \times \text{LTW}}{1 + (\delta R)^{\sigma} \times \text{LTW}}. \quad (3.11)$$

*Proof.* The proof is in [Supplementary Appendix B](#).  $\parallel$

Besides the structural parameters  $\sigma$ ,  $\delta$ , and the interest rate  $R$ , the reduced-form elasticity depends on a ratio we have defined as LTW. To get an intuitive sense of this ratio, consider a household

19. Besides very small  $\sigma$ s 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 address in Section 3.2.

whose only initial wealth is (the net worth of) housing and who has no current income. In that case,  $W_0 = (1 - \lambda_0)P_0H$  and the ratio  $\text{LTW} = \frac{\lambda_0 P_0 H}{y_1 + (1-d)P_1 H}$  represents a loan to future wealth ratio, with future wealth incorporating both human and financial (housing) wealth. For brevity, we refer to this ratio as a loan-to-wealth ratio. This ratio is endogenous and will in general differ substantially across households.<sup>20</sup>

Supplementary Appendix Figure A.5 illustrates the mapping between the EIS and the reduced-form elasticity under different LTW ratios. We see that the reduced-form elasticity can vary greatly for a given EIS, depending on LTW. Conversely, given an estimate of the reduced-form elasticity, there is very large variation in the EIS parameters that could be consistent with that estimate. This makes it difficult to infer the likely magnitude of the EIS from reduced-form evidence.<sup>21</sup>

### 3.2. Optimization frictions

The model presented above assumes that there are no optimization frictions (such as inattention or misperception). However, some households may be prevented from bunching due to such frictions, in which case our estimate of  $\sigma$  would be downward biased. To deal with this general problem in empirical research, Kleven and Waseem (2013) developed a non-parametric frictions adjustment based on the presence of strictly dominated regions of behaviour above notches. In their setting, strictly dominated regions above income tax notches were used to estimate the fraction of non-optimizing agents, 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 and 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 immediately above an LTV notch implies a large drop in future consumption, but allows for (slightly) 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 at all ( $\delta > 0$ ), there exists no non-negative EIS ( $\sigma \geq 0$ ) 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, we can structurally derive an LTV range above the notch that is inconsistent with any  $\sigma \geq 0$ . This range can be characterized as follows:

**Proposition 3 (Dominated region)** Under any  $\sigma \geq 0$ , choosing an LTV at the notch point  $\lambda^*$  dominates any interior LTV  $\lambda > \lambda^*$  for households whose counterfactual LTV satisfies

$$\lambda^* + \Delta\lambda \in \left( \lambda^*, \left( 1 + \frac{\Delta R}{R+1} \right) \lambda^* \right). \quad (3.12)$$

*Proof.* The proof is in Supplementary Appendix B. ||

20. The result in Proposition 2 also implies that, when  $\sigma$  converges to zero, the value of  $\varepsilon$  converges to  $R/(1+R)$ . This is a lower bound on the reduced-form elasticity and represents the wealth effect. This particular result is driven by the two-period assumption. In Supplementary Appendix C, we generalize our results to a multi-period version of the baseline model (retaining the other simplifying assumptions of this model) and show that the pure wealth effect (lower bound on  $\varepsilon$ ) is smaller in this case.

21. In the numerical example presented in the figure (in which  $\delta = R = 1$ ), the reduced-form elasticity is bounded from below by  $R/(1+R) = 0.5$  corresponding to the pure wealth effect (see equation (3.11) for  $\sigma = 0$ ).

We estimate the fraction of non-optimizers as the observed density mass in proportion to counterfactual density mass within the dominated region defined in (3.12). We assume that the fraction of non-optimizers in the dominated region is a good proxy for optimization frictions elsewhere in the distribution (where the amount of friction cannot be observed). Denoting the fraction of non-optimizers by  $a$ , the friction-adjusted bunching mass equals  $\frac{B}{1-a}$  and the frictionless LTV response equals  $\frac{\Delta\lambda}{1-a}$ . These are the reduced-form statistics that enter into the structural estimation of  $\sigma$ .

It is worth noting that the dominated regions from which we estimate  $a$  are very small. As shown in Figure 3, the largest notch is at  $\lambda^* = 80\%$  where the interest jumps by  $\Delta R = 0.5\text{pp} = 0.005$ . Using (3.12), this gives a dominated LTV range of approximately  $(80\%, 80.2\%)$  for  $R \simeq 1$ . For example, someone located at an LTV of  $80.1\%$  and with a house worth £200,000 (the average house value around this notch) would have to inject only £200 in order to get the 0.5pp reduction in the annual interest rate on the entire mortgage (worth roughly £3,000 in lower interest payments). If the household does not take this investment opportunity, we attribute it to optimization friction.

The strengths and weaknesses of this approach were discussed in detail in Kleven and Waseem (2013) and Kleven (2016). The key strength of the approach is that it provides a direct empirical measure of optimization friction that does not rely on any parametric assumptions on the specific structure and distribution of frictions, and in fact is completely agnostic about the sources of friction. Alternative approaches in the literature either ignore optimization friction or rely on strong structural assumptions about such frictions. The limitation of the approach is the assumption that the fraction of non-optimizers in the dominated region (where the incentive to respond is the strongest) equals the fraction of non-optimizers above the dominated region. As an example, consider the case where the optimization friction takes the form of a fixed adjustment cost. As discussed by Kleven and Waseem (2013), for a given distribution of adjustment costs, the fraction of non-optimizers  $a$  is in general increasing in the distance to the notch as the utility gain from bunching is declining. The value of  $a$  obtained from the dominated region is therefore a *lower bound* on the average fraction of non-optimizers above the notch. The magnitude of the bias depends on the distribution of adjustment costs. In the special case where a fraction of agents have zero adjustment costs (“optimizers”) and a fraction have prohibitively large adjustment costs (“non-optimizers”), the parameter  $a$  accurately captures the fraction of non-optimizers and yields unbiased estimates.

As an alternative, Kleven and Waseem (2013) also propose an *upper bound* on the amount of friction. This approach assumes that all density mass in the hole—not just the fraction estimated from the narrower dominated region—can be explained by friction rather than by heterogeneity in the true structural elasticity. Under this assumption, the true frictionless density distribution would look like Figure 5B, even if the observed density features a gradual hole. If the notch is not too large (so that we have  $\lambda^I \approx \lambda^* + \Delta\lambda$ ), the behavioural response  $\Delta\lambda$  can be estimated as the point of convergence between the observed and counterfactual distributions.<sup>22</sup> This point corresponds to that which is obtained from  $\frac{B}{1-\bar{a}}$ , where  $\bar{a}$  is measured using all density mass on the segment  $(\lambda^*, \lambda^* + \Delta\lambda)$ . Using the two approaches outlined above, it is possible to bound the amount of friction and therefore the true structural EIS.

22. For reasons explained in Kleven and Waseem (2013) and Kleven (2016), the assumption that  $\lambda^I \approx \lambda^* + \Delta\lambda$  is reasonable in most, if not all, notch applications.

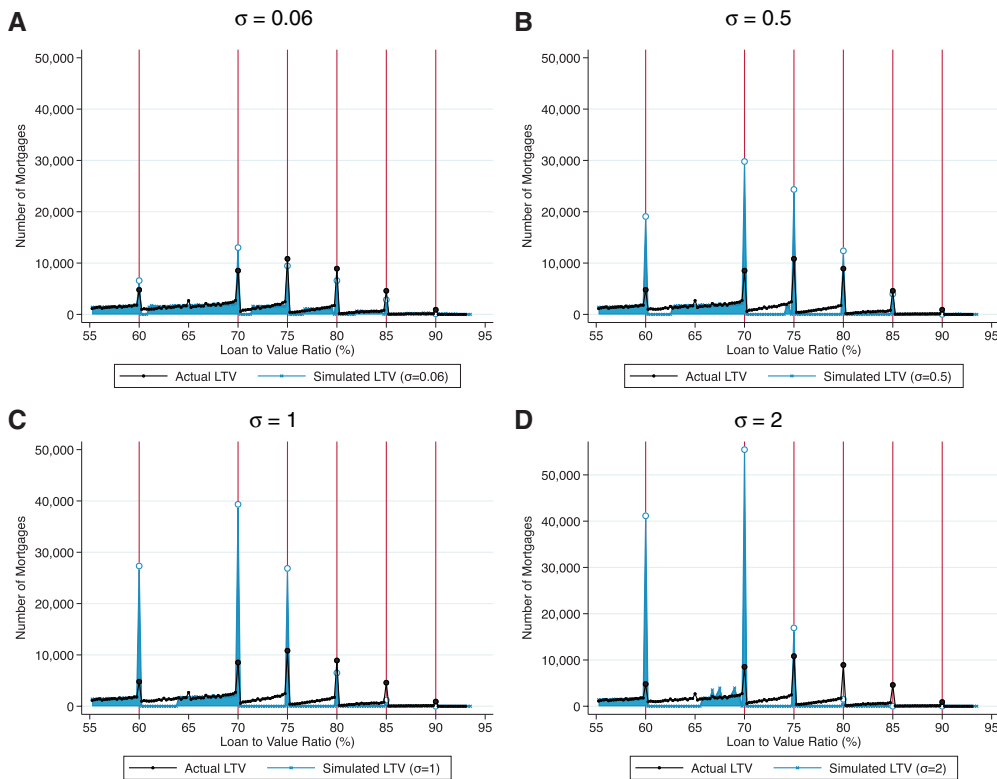


FIGURE 6

Observed versus simulated LTV distributions when varying the EIS.

*Notes:* The figure shows simulations of a model introduced in Section 3 for a range of EIS values. The lighter lines show the predicted LTV distribution if households choose leverage optimally according to the model. The black lines show the empirical LTV distribution. The upper left hand corner has  $\sigma = 0.06$ , which is the EIS that minimizes the MSE of the predicted bunching masses. Higher EIS values predict far greater bunching masses than found in the data, with a large share of households jumping more than one notch in the LTV distribution to exploit lower interest charges. The distribution largely hollows out between notches, in contrast to the data.

### 3.3. Identification of the EIS: numerical simulations

As discussed above, it is not immediately apparent how the EIS can be identified from bunching, because the estimating indifference equation (3.8) contains other parameters: the discount factor, future house prices, and future income. In this section, we present simulations of the global LTV distribution under different parameter configurations, which illustrate that *only* the EIS can be used to fit the observed distribution. While other parameters play some role, their impacts on bunching responses are very minor.

Figure 6 compares the observed LTV distribution to simulated LTV distributions under four different EIS scenarios. The other parameters of the model are assigned reasonable values that do not vary across the different EIS scenarios.<sup>23</sup> The distribution of initial wealth  $W_0$  is calibrated using equation (3.5) in order to replicate the counterfactual LTV distribution shown in Figure 4. In this counterfactual scenario, we assume that each borrower faces a flat interest rate  $R$  given by

23. Specifically, the discount factor is set at an annual rate of  $\delta = 0.96$  (a common value in the literature), real house price growth is set at an annual rate of  $P_1/P_0 = 1.026$  (the historical average in the U.K.), the depreciation rate is set at  $d = 0.025$  (taken from the literature), while for simplicity real income is assumed to be constant over time  $y_1 = y_0$ .

the observed rate at the counterfactual location.<sup>24</sup> Having calibrated the model in this way, the simulated LTV distribution is based on introducing the notched interest rate schedule and letting borrowers choose LTV optimally by comparing utility levels at their best interior location and at the five notch points.

The resulting distributions under each  $\sigma$  is shown in the four panels of Figure 6. Figure 6A sets  $\sigma = 0.06$ , corresponding to the EIS that minimizes the mean squared error (MSE) between the simulated bunching masses and the observed bunching masses. Figure 6B–D consider values of the EIS commonly used in the literature ( $\sigma = 0.5, 1, 2$ ).<sup>25</sup> The figure shows clearly why the data “demands” a low EIS. With higher elasticities, households grossly over-respond to the interest rate notches and often skip several notches in search of lower rates. As a result, the simulated distributions have almost no mass above the 70% LTV threshold, except at the notch points. There is *too little* bunching at the highest notches (as borrowers tend to skip these notches) and *too much* bunching at the lowest notches. This contrasts with the simulated distribution under  $\sigma = 0.06$ , which does a far better job of matching the data.<sup>26</sup>

While these results show that a low EIS can reconcile model and data, [Supplementary Appendix Figure A.6](#) shows that a low EIS is the *only* way of reconciling the two. The top panel repeats the best fit from the previous figure, *i.e.*  $\sigma = 0.06$  with the other parameters set to realistic values. The lower panel instead sets  $\sigma = 1$  and calibrates all the other parameters (discount rate, house price growth, and income growth) so as to minimize the MSE of the simulated bunching masses. Even when all the other parameters are fine-tuned to satisfy this single objective, the model provides a very poor fit to the data. This is because bunching responses are relatively insensitive to these other parameters.<sup>27</sup> What is more, calibrating the remaining parameters in this way leads to highly unrealistic values, including an annual discount rate of  $\delta = 0.24$  and annual income growth of  $-60\%$ .

The only real threat to identification is the presence of optimization frictions such as inattention, inertia, or misperception. While  $\sigma$  is the only parameter that can close the gap between model and data within a frictionless model, a sufficient amount of friction would be another way of closing the gap. It is therefore crucial that we have an empirical handle on the amount of friction from the dominated regions. Using the approach laid out in the previous section, [Supplementary Appendix Figure A.7](#) explores if there is enough optimization friction to justify a much higher value of the EIS. The figure is constructed like Figure 6, but the simulated LTV distributions have been adjusted to account for the presence of non-optimizers. Specifically, denoting by  $a_n$  the fraction of households observed in the dominated region above notch  $n$ , the simulations assume that the fraction  $a_n$  of households between the notch  $n$  and the next notch  $n + 1$  are stuck at their counterfactual LTV. This exercise creates less bunching at the notches and greater mass between the notches, but it does not fundamentally alter our conclusions. The EIS that provides the best fit is still small ( $\sigma = 0.12$ ) and standard EIS values provide very poor fits. Therefore, while it is clearly important to account for optimization frictions, the dominated regions suggest that there is not nearly enough friction in our setting to justify large structural elasticities.<sup>28</sup>

24. That is, borrowers get individualized flat interest rates based on their counterfactual bracket location. Our results are insensitive to the specific assumption we make about the counterfactual interest rate level.

25. Panel C ( $\sigma = 1$ ) repeats the simulation shown in Figure 2 discussed earlier.

26. The fit is quite impressive when considering the crudeness of assuming a single  $\sigma$  throughout the LTV distribution. This assumption is relaxed in our local estimations presented below.

27. We show in Section 5 that this remains true in a rich stochastic lifecycle model where several other parameters can be changed.

28. The model can fit the data better if we allow for heterogeneity in  $\sigma$ . Specifically, this helps with matching the gradual recovery of the LTV distribution—as opposed to the sharp holes—above notches. We have conducted such an



The simulation exercises presented here are useful for illustrating how bunching can identify the EIS and why the data call for a small elasticity. However, these simulations do not yield precise estimates of the EIS, because they are based on fitting global distributions assuming that a single elasticity applies everywhere. The next section presents local bunching estimations (using the estimating indifference equation characterized above) in which we relax this assumption.

#### 4. ESTIMATING THE EIS: SIMPLE MODEL

##### 4.1. *Bunching estimation*

In this section, we use bunching to estimate the EIS and the reduced-form borrowing elasticity based on the simple framework developed above. The next section extends the analysis to our full structural model.

We first consider all notches together by pooling the data into a single average notch. For each notch point  $n$  and each mortgage  $i$ , we calculate a normalized LTV as  $LTV_{in} = LTV_i - n$ .<sup>29</sup> 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 7 in which the actual LTV distribution is the series in dots and the counterfactual LTV distribution is the series in crosses, with their frequencies in 0.25pp bins given on the left axis. The counterfactual is obtained using the panel method described in Section 2.4, with the global distribution shown in Figure 4. The figure also plots the pooled conditional interest rate at each LTV with green squares (right axis) obtained from the non-parametric regression described in Section 2.3.<sup>30</sup> The graphs are shown with 95% confidence bands computed by bootstrapping.<sup>31</sup>

The figure displays the two key empirical moments that we will use later: the interest rate jump at the notch,  $\Delta r$ , and the amount of bunching scaled by the counterfactual density at the threshold,  $b \equiv B / [g_0(\lambda^*) \times \text{binwidth}]$ .<sup>32</sup> As shown in equation (3.10), this bunching statistic is approximately proportional to the local average LTV response. In the figure, we distinguish between two different bunching statistics:  $b(\text{raw})$  and  $b$ . The first estimate is based on the raw data shown in the figure, while the second (smaller) estimate adjusts for the presence of round-number bunching. As discussed earlier, we adjust for round-number bunching by exploiting the fact that some banks do not feature certain notches during certain periods.<sup>33</sup>

exercise, assuming that  $\sigma$  is independently distributed. The distribution of  $\sigma$  that provides the best fit is very tight around the point estimates reported above, *i.e.* we estimate very little heterogeneity in  $\sigma$ . The model would be able to fit the data even better if we allow for a joint distribution of  $\sigma$  and LTV  $\lambda$ . However, such simulations are computationally demanding and ultimately have a relatively low pay-off. Obtaining a precise fit of the entire hole does not improve upon the estimation of the *average*  $\sigma$  as compared to the simpler local bunching estimation presented below. Fitting the hole does give the entire distribution of  $\sigma$ , but this relies heavily on the modelling assumptions (including assumptions about the distribution of frictions).

29. 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.

30. The pooled interest rate line is a weighted average of conditional interest rates from the regression (2.1) implemented locally around each notch.

31. The bootstrap procedure draws 100 samples of mortgages from the observed sample, with replacement and stratifying by notch.

32. We measure bunching  $B$  as excess mass in the actual distribution relative to the counterfactual distribution in a 1pp range below the threshold, thus allowing for some overshooting by bunching households. Because bunching is extremely sharp in our setting, the results are very robust to reducing the width of this range.

33. [Supplementary Appendix Figure A.8](#) shows the LTV distribution in the no-notch subsample, *i.e.* mortgages located close to one of the five thresholds in a bank-by-time cell that does not feature an interest notch at that threshold. Panel A of the figure shows that the no-notch subsample is a very small fraction of the full sample. Panel B of the figure zooms in on the no-notch subsample, demonstrating that there is a modest amount of round-number bunching. To adjust






FIGURE 7

## Bunching estimation when pooling notches

*Notes:* The figure shows the actual ( $f(\lambda)$ ) with black dots) and counterfactual ( $f_0(\lambda)$  with crosses) distributions of LTV ( $\lambda$ ), pooling all notches (60%, 70%, 75%, 80%, and 85%). The squares show the conditional interest rate in each LTV bin from the regression described in Section 2.3 and the footnote to Figure 3. The counterfactual LTV distribution is obtained using the method outlined in Section 2.4 and the footnote to Figure 4. The figure also shows the jump in the conditional interest rate at the notch  $\Delta r$ ; the (raw) normalized amount of bunching in the actual distribution  $b(\text{raw})$ , calculated as described in Section 3; and  $b$  which is the amount of bunching net of round number bunching.

The following findings emerge from Figure 7. First, there is large and sharp bunching equal to five times the height of the counterfactual distribution at the notch, or about 10% less when accounting for round-number bunching. Second, there is a clear gap between the actual and counterfactual distributions to the right of the notch. This is the “hole” in which the bunching households would have been observed absent the notched interest rate schedule. The hole extends to around 3.5pp above the notch, implying that the *most* responsive households reduce their LTV by 3.5pp in response to the average interest rate notch of 0.25pp. This upper-bound response is well below the next notch (and the confidence band is tight), validating our assumption in Proposition 1 that bunchers move only one notch.

Third, when comparing the actual and counterfactual densities immediately above the notch, we see that about 30% of borrowers are stuck in a dominated region.<sup>34</sup> As discussed in Section 3.2, we interpret the observed mass in the dominated region as a reflection of optimization frictions such as switching costs, inattention, and misperception. In this setting, there is an additional reason for locating just above notches: the fact that some borrowers have mortgages in banks that do not feature that particular notch at the time of their loan origination. However, the impact of no-notch banks on the hole is limited by the fact that they represent a very small fraction of the data as mentioned above (see [Supplementary Appendix Figure A.8](#)). Moreover, locating just

for round-number bunching, we estimate the degree of bunching at the notches using the no-notch subsample. As is common in the bunching literature, we estimate the counterfactual LTV distribution using a 5th order polynomial fitted to the data collapsed into bins with a width of 0.25 percentage points and excluding data within 1% of the notch. We then subtract the estimated normalized bunching in the no-notch sample,  $b_{RN}$ , from the bunching moment of interest,  $b_{\text{raw}}$ , to calculate our estimate of the bunching due to the interest rate notches.

34. The dominated regions (as characterized in Proposition 3) are even smaller than the 0.25pp LTV bins depicted in the figure, but we show later that 30% is nevertheless about right.

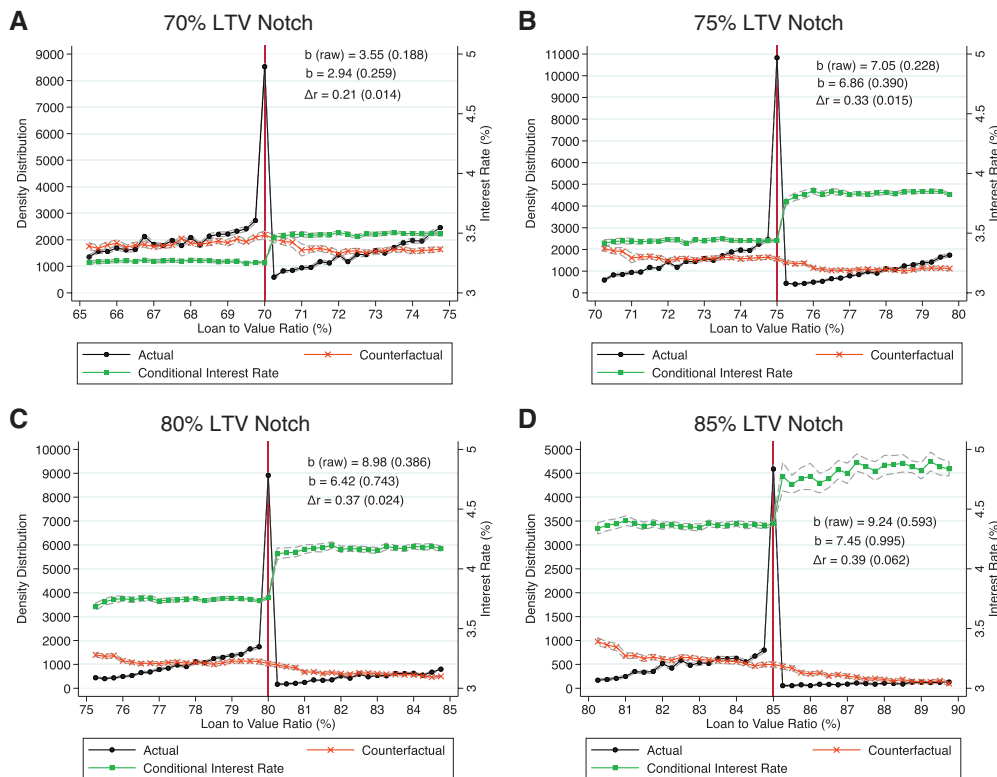


FIGURE 8

Bunching estimation at individual notches.

*Notes:* The figure shows the actual ( $f(\lambda)$ ) with black dots and counterfactual ( $f_0(\lambda)$ ) with crosses distributions of LTV ( $\lambda$ ), at the notches at 70% (Panel A), 75% (Panel B), 80% (Panel C) and 85% (Panel D) LTV. The squares show the conditional interest rate in each LTV bin from regression described in Section 2.3 and the footnote to Figure 3. The counterfactual LTV distribution is obtained using the method outlined in Section 2.4 and the footnote to Figure 4. The figure also shows the jump in the conditional interest rate at the notch  $\Delta r$ ; the (raw) normalized amount of bunching in the actual distribution  $b(\text{raw})$ , calculated as described in Section 3; and  $b$  which is the amount of bunching net of round number bunching.

above a threshold in a no-notch bank may in fact be viewed as an optimization friction in the form of bank switching costs. As [Supplementary Appendix Figure A.9](#) shows, borrowers located above a threshold in a no-notch banks could typically get a large discount by switching to a similar product in a different bank and moving below the threshold, and their unwillingness to do so must be related either to a friction in bank choice or to unobserved services provided by the no-notch bank. When we adjust for optimization frictions using the dominated regions, we assume either that (i) all mass in the dominated region is due to friction or (ii) that only the mass coming from notched banks is due to friction.

Figure 8 shows bunching evidence for individual notches, but is otherwise constructed in the same way as the previous figure. The interest rate jumps shown in this figure are somewhat smaller than those reported in Figure 3, because the interest rate jumps shown here apply to the refinancer sample (as opposed to the full population of people with mortgages) and therefore to a different composition of mortgages. The evidence from the individual notches is qualitatively consistent with the evidence from the pooled notch. The amount of bunching  $b$  is increasing in the size of

TABLE 2  
*From bunching to the EIS: simple model*

Statistic	Notch					
	60	70	75	80	85	Pooled
Panel A: Bunching evidence						
$r(\%)$	3.17 (0.01)	3.25 (0.00)	3.44 (0.00)	3.76 (0.00)	4.38 (0.01)	3.42 (0.01)
$\Delta r(\%)$	0.10 (0.01)	0.21 (0.01)	0.33 (0.02)	0.37 (0.02)	0.39 (0.06)	0.25 (0.01)
$b$	0.96 (0.17)	2.94 (0.26)	6.86 (0.39)	6.42 (0.74)	7.45 (0.99)	4.45 (0.20)
$a$	0.58 (0.05)	0.21 (0.02)	0.30 (0.03)	0.15 (0.02)	0.08 (0.03)	0.29 (0.01)
$b_{\text{Adj}}$	2.31 (0.49)	3.73 (0.35)	9.87 (0.60)	7.59 (0.89)	8.11 (1.16)	6.30 (0.30)
$\Delta\lambda_{\text{Adj}}$	0.67 (0.14)	1.06 (0.09)	3.32 (0.18)	2.68 (0.32)	3.71 (0.70)	1.93 (0.09)
Panel B: Elasticities						
EIS $\sigma$	0.03 (0.01)	0.03 (0.00)	0.17 (0.02)	0.08 (0.02)	0.13 (0.05)	0.07 (0.01)
Reduced-form $\varepsilon$	0.53 (0.01)	0.53 (0.00)	0.60 (0.01)	0.56 (0.01)	0.58 (0.02)	0.55 (0.00)

*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,  $\Delta r$  is the interest rate jump at the notch, estimated as described in Section 2.3,  $b_{\text{Adj}}$  is our normalized bunching estimate, net of round number bunching and adjusted for optimization frictions, and  $\Delta\lambda_{\text{Adj}}$  is our estimate of the leverage response, estimated as described in Section 3. The EIS  $\sigma$  is estimated using a simple two-period model described in Section 3: the solution to the indifference equation (3.8). The reduced-form elasticity is obtained using (3.11), derived in Section 3. The standard errors, shown in parentheses, are obtained by bootstrapping the estimation routine, stratifying by notch, 100 times.

the interest rate jump, and the amount of mass just above the notch (friction) is decreasing in the size of the jump, exactly as one would expect.

#### 4.2. Elasticity estimation

We turn to the estimation of the EIS in Table 2. The table shows results for the five individual notches and for the pooled notch. Panel A starts by summarizing the statistics presented so far: the interest rate level below each notch  $r$ , the interest rate jump  $\Delta r$ , and the bunching statistic  $b$ . The panel also shows the fraction of non-optimizers  $a$  obtained from the dominated region, the friction-adjusted bunching statistic  $b_{\text{Adj}} = b/(1-a)$ , and the implied frictionless LTV response  $\Delta\lambda_{\text{Adj}}$ . The average LTV response using the pooled notch is close to 2pp, *i.e.* households are willing to reduce borrowing by an average of 2% of their house value in order to avoid the interest jump of 0.25pp. Apart from the 80% notch, the LTV response is monotonically increasing in the size of the interest jump, from 0.7pp at the lowest notch to 3.7pp at the highest notch. Given the notches are at least 5pp apart, the estimated LTV responses are consistent with the assumption (made in the theoretical model) that bunchers move only one notch down.

Panel B of the table turns the reduced-form evidence into elasticities using the simple structural model from Section 3. The EIS  $\sigma$  is based on the estimating indifference equation shown in Proposition 1, and the reduced-form borrowing elasticity  $\varepsilon$  is based on the equation on shown in Proposition 2. Besides the bunching moments, the estimation of elasticities requires us to set a

few additional parameters that are not directly observed: the discount factor  $\delta$ , future income  $y_1$ , and house price growth net of depreciation  $\Pi_1 = (1 - d) \frac{P_1}{P_0}$ . As discussed in detail in the previous section, the estimation is not sensitive to the assumptions we make about these parameters.<sup>35</sup> Arguably, the crudest assumption in this exercise is the two-period nature of the model, but we have repeated the exercise for a multi-period version of the model and the EIS estimates are virtually identical (the multi-period extension is presented in [Supplementary Appendix C](#)).<sup>36</sup>

As shown in the table, the EIS is small. It ranges from 0.03 to 0.17 across the different notches, and the average elasticity obtained from the pooled notch is close to 0.1. The pooled estimate is very close to the calibrated EIS obtained from the global simulation exercise in the previous section. Translating the structural EIS estimates into reduced-form borrowing elasticities, we obtain stable values of  $\varepsilon$  across the different notches, all in the neighbourhood of 0.5. As implied by the result in Proposition 2, when  $\sigma$  converges to zero, the value of  $\varepsilon$  converges to  $R/(1+R) \simeq 0.5$  and represents a pure wealth effect. Given the low values of the EIS deriving from our structural model, most of the borrowing response to interest rates is due to the wealth effect.

The elasticity estimates in Table 2 are corrected for optimization frictions using the Kleven–Waseem adjustment factor  $a$ , *i.e.* the fraction of borrowers who are stuck in the dominated regions above notches. The calculation of the dominated regions was based on the *average* interest jumps at notches in an empirical setting that features heterogeneity in the size of interest jumps across banks and products. In particular, because some banks do not feature certain notches during certain time periods, it is conceivable that the mass immediately above notches disproportionately comes from such no-notch contracts. In this case, we would overstate the amount of friction and therefore the EIS.<sup>37</sup> On the other hand, as discussed in section 3.2, there are other reasons why the Kleven–Waseem friction adjustment may understate the amount of friction. To investigate the sensitivity to these issues, Table 3 puts bounds on the friction adjustment and the EIS by considering the following scenarios: (i) no friction adjustment, (ii) a friction adjustment based on the dominated region and notched banks only, (iii) a friction adjustment based on the dominated region and all banks (our baseline estimates), and (iv) a friction adjustment assuming that all mass in the hole (and not just mass in the dominated region) is due to friction. As implied by the discussion in Section 3.2, scenarios (i) and (ii) provide *lower bounds* on the amount of friction and therefore on the EIS, while scenario (iv) provides an *upper bound* on the amount of friction and the EIS. The table shows that the EIS estimates are relatively tight. Even under the extreme assumption that all mass in the hole is due to friction, the average EIS is fairly modest at around 0.3.

The results in Table 2 are estimates of the average EIS in the population, but there may be subgroups with a higher EIS. In fact, an important advantage of our rich data is the ability to study heterogeneity in the EIS across subsamples. Hence, in Table 4 we explore heterogeneity in the EIS along a number of dimensions. For each covariate, we partition the sample into four quartiles and separately estimate the EIS in each quartile. The table reveals only modest heterogeneity in the EIS. The elasticity is larger for older households, those with lower income, those with higher

35. We assume that the annual discount factor equals  $\delta = 0.96$ , that future income equals current income ( $y_1 = y_0$ ), and that annual house price growth equals its historical average in the U.K. ( $P_1/P_0 = 1.026$ ) net of a depreciation rate taken from the literature ( $d = 0.025$ ). The period length is set equal to 3.34 years, corresponding to the average time between refinance events in the data.

36. While the structural EIS is not sensitive to the number of periods, the reduced-form borrowing elasticity does change when changing the number of periods (all else being equal). Adding periods to the model tends to lower the reduced-form elasticity by lowering (spreading out) the wealth effect of the interest rate.

37. However, the potential bias is limited by the fact that the fraction of no-notch contracts is very small as we saw in [Supplementary Appendix Figure A8](#).

TABLE 3  
Bounding optimization frictions and the EIS

	Notch				
	70	75	80	85	Pooled
Panel A: Adjustment factor $a$					
(1) Notched banks only	0.11 (0.02)	0.15 (0.02)	0.15 (0.03)	0.03 (0.01)	0.12 (0.01)
(2) Dominated region	0.21 (0.02)	0.30 (0.03)	0.15 (0.02)	0.08 (0.03)	0.22 (0.01)
(3) Entire hole	0.67 (0.05)	0.60 (0.02)	0.57 (0.04)	0.40 (0.09)	0.61 (0.02)
Panel B: Elasticity of intertemporal substitution $\sigma$					
(4) Unadjusted	0.02 (0.00)	0.08 (0.01)	0.06 (0.01)	0.11 (0.04)	0.05 (0.01)
(5) Dominated region: notched banks only	0.02 (0.00)	0.11 (0.01)	0.08 (0.02)	0.11 (0.04)	0.07 (0.01)
(6) Dominated region: all banks	0.03 (0.00)	0.17 (0.02)	0.08 (0.02)	0.13 (0.05)	0.09 (0.01)
(7) All mass in the hole is friction	0.16 (0.05)	0.50 (0.07)	0.31 (0.08)	0.30 (8.53)	0.30 (0.03)

*Notes:* The table shows how the estimated EIS is affected by assumptions on optimization frictions. The top panel of the table shows the friction adjustment factor  $a$  estimated in three different cases. Row (1) shows the friction adjustment based on mass in the dominated region using only notched banks, row (2) shows the friction adjustment based on mass the dominated region using all banks (our baseline estimates), while row (3) shows the friction adjustment assuming that all mass in the hole is due to friction. The bottom panel of the table shows the estimated EIS when not adjusting for optimization friction (in row (4)), and when adjusting for friction using each of the three measures provided in the top panel (in rows (5)–(7)). As explained in the main text of the article, the EIS estimates provided in rows (4) or (5) are in general lower bounds, whereas the EIS estimate provided in row (7) is an upper bound. The upper bound is based on the extreme assumption that all density mass in the hole—not just the mass in the much narrower dominated region—can be explained by friction rather than by heterogeneity in true preferences (*i.e.* true preferences are assumed to be homogeneous in the population).

TABLE 4  
Heterogeneity in the EIS

Covariate	Quartile			
	1	2	3	4
Age	0.05 (0.01)	0.09 (0.02)	0.10 (0.02)	0.15 (0.08)
Household income	0.09 (0.02)	0.08 (0.01)	0.07 (0.01)	0.05 (0.01)
Loan-to-income	0.02 (0.01)	0.05 (0.01)	0.08 (0.01)	0.07 (0.02)
Income growth	0.05 (0.01)	0.06 (0.02)	0.07 (0.01)	0.07 (0.02)
House price growth rate	0.06 (0.02)	0.05 (0.01)	0.04 (0.01)	0.13 (0.03)
Interest rate change (passive)	0.02 (0.01)	0.06 (0.02)	0.11 (0.03)	0.11 (0.03)

*Notes:* The table shows the heterogeneity in our estimated EIS  $\sigma$  (using the pooled average notch) by age, income, loan to income (LTI), income growth, house price growth, interest rate change since the previous mortgage (assuming passive borrower behaviour). For each covariate, we partition the refinancer panel into four quartiles and separately estimate  $\sigma$  in each quartile. The standard errors, shown in parentheses, are obtained by bootstrapping the estimation routine, stratifying by notch, 100 times.

LTI ratios, and those experiencing the fastest house price growth.<sup>38</sup> However, the differences are only modest, with the EIS ranging from 0.02 to only 0.15 across all the subgroups considered.

To conclude, we have translated bunching at mortgage notches into structural EIS parameters through the lens of the simplest possible dynamic model. In this model, the only way households can reallocate consumption over time is through borrowing against their house. This implies that bunching at mortgage notches maps directly into intertemporal consumption reallocation and therefore reveal the EIS. By assuming that households can bunch only by reducing current consumption, as opposed to drawing down liquid assets, we are estimating an *upper bound* on the EIS. The presence of liquid assets gives households an extra margin of adjustment and will in general add to bunching, unless liquidity demand is completely inelastic to interest rates.<sup>39</sup> That is, a given amount of observed bunching is consistent with a smaller consumption reallocation in the presence of liquidity/portfolio effects, and by implication a smaller value of the EIS. This makes our small EIS estimates very informative.

As implied by this discussion, exactly identifying the EIS requires a model of how liquidity demand responds to the interest rate. This calls for a model that allows for uncertainty and a precautionary savings motive, in which case the responsiveness of liquidity depends on the amount of uncertainty and risk preferences. In our baseline model with Constant Relative Risk Aversion (CRRA) preferences, the coefficient of relative risk aversion is the inverse of the EIS, and so a low EIS will make households very risk averse and therefore make liquidity demand inelastic. To allow for more flexibility in liquidity responses, it is natural to separate the EIS and risk aversion preferences using Epstein–Zin preferences. In the following section, we develop a stochastic lifecycle model that allows for all of these features in order to demonstrate how the EIS relates to bunching in a richer and more realistic setting. However, for the high-level reasons discussed here, the liquidity channel does not give rise to larger estimates of the EIS, all else equal.<sup>40</sup>

## 5. ESTIMATING THE EIS: FULL STRUCTURAL MODEL

### 5.1. *Model*

We now turn to estimating the EIS in our extended structural model. Households live for  $T$  periods (years) and choose consumption, housing, liquidity, mortgage debt, and bequests. Future house prices and income are uncertain. The mortgage interest rate is a step function of LTV, corresponding to the notched interest schedule in the U.K. Households may buy and sell housing, and they may hold an additional liquid asset. As noted above, liquidity demand may be confounded with consumption demand in driving bunching behaviour. The presence of a liquid asset in the model allows a quantitative assessment of the relative importance of these two channels in driving bunching. We present the basic structure of the model here and relegate its full analysis to [Supplementary Appendix D](#).

38. This analysis considers each dimension of heterogeneity in isolation, raising the concern that if the covariates are correlated, it is unclear which covariate to attribute the heterogeneity to. To address this, in results available from the authors, we split the data at the median of age, income, and LTI ratio, calculated the EIS separately in these eight subgroups, and then regressed the estimated EIS on the three covariates. The results confirm the patterns in Table 4.

39. Adding fixed (price inelastic) liquid wealth to our model would have no impact on our results, because in that case borrowers still bunch using only consumption.

40. When moving to the full structural model below, we generalize in more dimensions than the liquidity channel. Due to the combination of generalizations, it is possible for the full structural estimates to be larger than the baseline estimates presented above.

**5.1.1. Setup.** Households have preferences over non-durable consumption  $c_t$  and housing services  $H_t$ , defined recursively as

$$V_t = \left( \left( c_t^\alpha H_{t+1}^{1-\alpha} \right)^{\frac{\sigma-1}{\sigma}} + \delta \left( E_t \left\{ V_{t+1}^{1-\gamma} \right\} \right)^{\frac{1}{1-\gamma} \frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (5.13)$$

where  $V_t$  is the net present value of utility at period  $t$ ,  $\sigma$  is the EIS,  $\gamma$  is the coefficient of relative risk aversion, and  $\alpha$  is the share of housing in overall consumption. These preferences follow Epstein and Zin (1989) and Weil (1990) and allow for a distinction between risk aversion and the EIS. This helps ensure that our EIS estimates are not influenced by the dual role this parameter plays in standard CRRA preferences. The unit elasticity of substitution between housing and non-housing consumption (Cobb–Douglas) is justified by the observation that expenditures on housing services have historically been a constant share of total household expenditures (see *e.g.* Piazzesi and Schneider, 2016).

Households value end-of life wealth  $W_{T+1}$  via a reduced-form bequest motive, *i.e.*

$$V_{T+1} = \Gamma W_{T+1}, \quad (5.14)$$

where  $\Gamma$  is a parameter governing the intensity of bequest motives.

Households enter period  $t$  with two assets and one liability. The assets are housing  $H_t$  with a market price of  $P_t$ , and a liquid asset  $L_t$  denominated in units of the numeraire consumption good. The liability is mortgage debt  $D_t$ , again denominated in units of the consumption good. The real gross mortgage interest rate is given by  $R_t = 1 + r_t$ , while the liquid asset obtains zero nominal return. For simplicity, we abstract from short-term credit so that households' liquidity constraint is given by  $L_t \geq 0$ .

To conserve on computational power, we assume that housing quality  $H_t$  can take on three values normalized to  $\{1, 1.2, 1.4\}$ . This is sufficient to allow for a lifecycle pattern featuring increasing housing in the beginning of life and decreasing housing at the end of life. With this discrete grid, moving costs turn out to be relatively unimportant. Our initial simulation exercises showed that moving costs occasionally delay moving by a year or two, but do not change the qualitative lifecycle pattern. We therefore abstract from the moving cost in what follows.

Mortgage contracts have fixed maturities of  $m$  years, after which a penalizing reset interest rate kicks in. We assume this reset rate is sufficiently high that households never refinance after this point. Households also face an early repayment penalty if they choose to refinance before the mortgage matures. Penalties equal to 5–10% of the outstanding loan are common in the U.K. Simulating the model with penalties of this magnitude shows that households virtually never refinance early. Accordingly, we assume that households refinance after  $m$  years unless they move (in which case the prepayment penalty is waived as is typically the case in the U.K.). We assume a simple amortization schedule with constant annual repayments ensuring full repayment by age 70. (The typical mortgage contract in the U.K. requires full repayment before the borrower reaches the age of 70.) That is, the amortization rate is given by

$$\mu_t = \frac{1}{70 - \text{Age} + 1}.$$

Of course, households have ample opportunities to readjust amortization by repaying or extracting equity when refinancing. We set the terminal period to  $T = 70$ . This understates average longevity in the U.K., but households are typically inactive in the mortgage market after this age. We may think of the bequest function  $V_{T+1}$  as capturing the overall preference for wealth at age 70 combining retirement and bequest motives.



When a household refinances, it must pay a fixed mortgage origination fee of  $\Omega$  consumption units. Houses depreciate in quality at a rate  $d$  in each period. For simplicity, we assume that households maintain their houses in each period so as to replace their depreciated value, *i.e.* pay a maintenance fee of  $d \cdot P_t H_t$  in period  $t$ . Households obtain an income stream of  $y_t$  consumption units in period  $t$ .

With these features, the household's budget constraint is given by

$$\begin{aligned} c_t = & y_t + (1 - \pi_t)L_t - L_{t+1} \\ & + P_t((1 - d)H_t - H_{t+1}) \\ & + D_{t+1} - R_t D_t - \Omega I_t^R, \end{aligned} \quad (5.15)$$

where  $I_t^R$  is an indicator equal to 1 if the household refinances ( $D_{t+1} \neq (1 - \mu_t)D_t$ ), and  $\pi_t$  gives the inflation rate for non-housing consumption.

The interest rate is a spread over a risk-free base rate. The spread is a function of LTV, corresponding to the notched mortgage interest schedule in the U.K. Formally,  $R_t = R_t^0 + \rho(\lambda_s)$  where  $R_t^0$  is the base rate and where the spread  $\rho(\cdot)$  is a step function of the LTV ratio  $\lambda_s$  at the time of mortgage origination  $s$ . The LTV ratio is defined as  $\lambda_s \equiv \frac{D_{s+1}}{P_s H_{s+1}}$ .

Households have rational expectations, are forward looking, and optimize in each period subject to the aforementioned adjustment costs. This may seem at odds with our assessment that a significant fraction of households face optimization frictions (reflected in our estimates of  $a$  in Tables 2 and 3). However, in our empirical estimation of the model, we will use bunching moments adjusted for optimization frictions, as outlined in Section 3. We can therefore restrict attention to fully optimizing agents, as these moments reflect the amount of bunching that would prevail if households faced no optimization frictions. This has the advantage that we do not have to take a position on the exact form of optimization frictions facing households in our theoretical framework.

The model is solved computationally, with the full details of our approach described in [Supplementary Appendix D](#). We solve for the value of the EIS using the bunching moments along with an indifference equation similar to (3.8), but using the value functions arising from the extended model. Parameter values other than the EIS are calibrated to match features of the data or taken from the existing literature. The calibration is summarized in [Supplementary Appendix Table A.1](#).

## 5.2. Results

Table 5 shows EIS estimates resulting from the extended model. Standard errors are obtained by block bootstrapping with replacement. The first rows of the table restate our estimates of bunching  $b$ , the fraction of non-optimizers  $a$ , bunching adjusted for optimization friction  $b_{\text{Adj}} = \frac{b}{1-a}$ , and the implied LTV response for optimizers  $\Delta \lambda_{\text{Adj}}$ . We report these estimates at each individual notch as well as at an “average” notch. The latter is obtained as a weighted average of the estimates at individual notches, with weights proportional to the number of borrowers around each notch in the counterfactual distribution.<sup>41</sup>

41. When using the full structural model, we cannot estimate the EIS at a pooled notch like we did in Section 4. In the simple model, we considered a single refinancing episode in which borrowers made bunching decisions around a single notch. In the extended model, we consider a setting with repeated refinancing episodes in which borrowers make bunching decisions at time  $t$  anticipating the full menu of five notches in future refinancing episodes.

TABLE 5  
*From bunching to the EIS: full structural model*

Statistic	Notch					Average
	60	70	75	80	85	
$b$	0.96 (0.17)	2.94 (0.26)	6.86 (0.39)	6.42 (0.74)	7.45 (0.99)	4.11 (0.19)
$a$	0.58 (0.05)	0.21 (0.02)	0.30 (0.03)	0.15 (0.02)	0.08 (0.03)	0.31 (0.01)
$b_{Adj}$	2.31 (0.49)	3.73 (0.35)	9.87 (0.60)	7.59 (0.89)	8.11 (1.16)	5.57 (0.26)
$\Delta\lambda_{Adj}$	0.67 (0.14)	1.06 (0.09)	3.32 (0.18)	2.68 (0.32)	3.71 (0.70)	1.88 (0.09)
EIS $\sigma$	0.05 (0.01)	0.04 (0.01)	0.11 (0.01)	0.11 (0.02)	0.28 (0.15)	0.08 (0.01)

*Notes:* The table shows the results of our structural estimation as described in Section 5.  $b$  is our normalized bunching estimate as described in Section 3 and the footnote to Figure 8.  $a$  is the adjustment factor for optimization frictions (the number of individuals observed in the dominated region 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  $b_{Adj}$ .  $\sigma$  is the EIS that solves the full model described in Section 5. The standard errors, shown in parentheses, are obtained by bootstrapping the estimation routine, stratifying by notch, 100 times.

In the last row of Table 5, we report our estimates of the EIS  $\sigma$  at each notch. The EIS is small and stable across notches, except at the 85% LTV notch where the elasticity is somewhat larger. The average EIS is slightly below 0.1. This is essentially identical to the estimate deriving from the simple two-period model shown in Table 2, despite all the bells and whistles of the full structural model. This suggests that the low EIS estimates reported earlier are not driven by the simplifying assumptions of the two-period model, but rather by the magnitude of bunching observed in the data.<sup>42</sup> It also confirms the identification arguments made in section 3.3.

Why does model specification matter so little? Despite the greater complexity of the model developed here compared to the two-period model in Section 3, solving for the EIS ultimately boils down to a similar indifference equation. There are two main differences between the models. First, the continuation value in the two-period model is simply the utility of second-period consumption, as opposed to the much more involved continuation value  $V_{t+1}$  in the full lifecycle model. This has little quantitative implication for EIS estimates. This is because most factors that affect the continuation value affect it by similar magnitudes when bunching at the notch and when locating in the interior, and therefore roughly cancel out from the two sides of the indifference equation (3.8). Rather, it is the curvature of the continuation value with respect to debt that governs bunching motivations and therefore EIS estimates. But the curvature of the value function  $V_t$  with respect to wealth (and therefore debt) is approximately equal to the EIS, as is often the case in dynamic consumption models. Hence the bunching decision is roughly the same in the full model as it was in the two-period model.

Second, our full structural model allows for a liquid asset and for liquidity choice, while no liquid asset was available in the two-period model. This is a more substantive difference, because

42. Why does the 85% notch yield a larger EIS? Our model implies substantial equity injection at the counterfactual at the time of bunching for the average household at this notch. This means that households are short of liquidity when attempting to bunch at the 85% notch and must forgo some current consumption for capital repayment even at the counterfactual. This makes bunching at the 85% notch particularly painful and justifies the small amount of bunching observed in the data even at a higher EIS approaching 0.3. This result illustrates how our model incorporates liquidity constraints in its estimation, but also that the EIS remains low even when incorporating severe liquidity constraints. We do not observe households' wealth directly, but if households at this notch have more wealth than our model suggests, our estimated EIS is overstated.

the liquid asset gives households an additional margin of adjustment when making the bunching decision. With liquid assets, households confronting the bunching decision can now forgo either consumption or liquidity to lower their LTV to the notch. In equilibrium, households equalize the marginal value of consumption with the marginal value of liquidity, so it is optimal to forgo a combination of the two. The moderate responses to interest rates implied by the observed bunching suggest that households are inelastic in their demand for consumption *and* liquidity. With reasonable values of risk aversion and uncertainty—the main parameters governing liquidity demand—the liquidity margin facilitates bunching, which heightens the “puzzle” of households’ small responses to borrowing rates. In contrast, the simple model without liquidity is equivalent to a model with a binding liquidity constraint or one with perfectly inelastic liquidity demand. Hence, the model with liquidity should give smaller EIS estimates all else equal. However, because the EIS estimates from the baseline model were already small, incorporating liquidity choice has only a small impact on the estimation.

We have shown that our EIS estimates are largely robust to first-order changes in modelling assumptions, *i.e.* moving from a very simplified two-period model to a full dynamic lifecycle model. In this section, we subject our model to an additional battery of tests to explore its robustness to specification and parameter values. Table 6 shows EIS estimates (at the average notch) across a range of parameterizations. In each row, the estimate in boldface corresponds to the baseline assumption considered above.

We start by exploring robustness to household impatience. The first row shows robustness to the discount factor  $\delta$ , while the second row allows for hyperbolic discounting and varies the coefficient of present bias  $\beta$ .<sup>43</sup> As noted earlier, the discount factor (or present bias) largely governs the *level* of borrowing, while the EIS governs the *responsiveness* of borrowing to interest rates. As a result, changes in the discount factor  $\delta$  or the present bias factor  $\beta$  have relatively little impact on the estimated EIS. Even under extreme present bias ( $\beta = 0.3$ ), the EIS estimate is only moderately larger than in the baseline.

We next turn to risk aversion, which governs liquidity preferences and may affect EIS estimates. We explore values of  $\gamma$  ranging from risk neutrality ( $\gamma = 1$ ) to high levels of risk aversion. The last column shows the CRRA case  $\gamma = \frac{1}{\sigma}$ , which implies extreme risk aversion due to the low range of  $\sigma$  in our estimates. Despite the potential role of liquidity in affecting our estimates, the EIS remains relatively small for the entire range of risk preferences.

The remaining robustness checks alter household expectations along a number of dimensions. First, we alter interest rate expectations. The baseline interest rate process was calibrated from the U.K. yield curve, which currently reflects expectations of interest rates near zero for an extended period. We see that our EIS estimates are similar if households expect higher interest rates (shifting the yield curve upwards by 1pp, 2pp, and 3pp). Second, we see that the estimated EIS is insensitive to expectations of (mean) house price growth and not particularly sensitive to uncertainty about house prices. We alter the expected lifecycle profile of income from static income expectations to a trajectory implying real income growth at 7% per annum. Again, the implications for the estimated EIS are small. Finally, we alter income uncertainty by either changing the probability of unemployment or altering the replacement rate for unemployed workers. In both cases, the EIS remains moderate.

43. For this latter exercise, we alter the model so that households maximize the following function in each period

$$\left( (c_t^\alpha H_{t+1}^{1-\alpha})^{\frac{\sigma-1}{\sigma}} + \beta \delta (E_t \{ V_{t+1}^{1-\gamma} \})^{\frac{1}{1-\gamma} \frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

with  $V_{t+1}$  defined as before.

TABLE 6  
Robustness of EIS estimates

(1)	Discount factor $\delta$		0.7	0.9	<b>0.96</b>	0.99
			0.13	0.12	<b>0.08</b>	0.12
			(0.015)	(0.015)	<b>(0.011)</b>	(0.013)
(2)	Present bias factor $\beta$		0.3	0.5	0.7	<b>1</b>
			0.17	0.14	0.11	<b>0.08</b>
			(0.026)	(0.019)	(0.015)	<b>(0.011)</b>
(3)	Risk aversion $\gamma$		0	1	<b>2</b>	CRRA
			0.11	0.12	<b>0.08</b>	0.15
			(0.011)	(0.012)	<b>(0.011)</b>	(0.013)
(4)	Future interest rates		<b>+0pp</b>	+1pp	+2pp	+3pp
			<b>0.08</b>	0.11	0.12	0.12
			<b>(0.011)</b>	(0.014)	(0.013)	(0.013)
(5)	House price trend		-0.6%	0	<b>0.6%</b>	6%
			0.10	0.10	<b>0.08</b>	0.10
			(0.014)	(0.014)	<b>(0.011)</b>	(0.014)
(6)	House price variance		0	0.004	<b>0.006</b>	0.008
			0.08	0.10	<b>0.08</b>	0.16
			(0.009)	(0.011)	<b>(0.011)</b>	(0.012)
(7)	Lifecycle income profile	Peak	£44K	<b>£46K</b>	£56K	£80K
		Slope	0%	<b>0.7%</b>	2.7%	6.5%
			0.12	<b>0.08</b>	0.09	0.08
(8)	Unemployment probability		3%	<b>5%</b>	7%	10%
			0.09	<b>0.08</b>	0.11	0.12
			(0.014)	<b>(0.011)</b>	(0.013)	(0.015)
(9)	Replacement rate		<b>60%</b>	80%	100%	
			<b>0.08</b>	0.13	0.12	
			<b>(0.011)</b>	(0.013)	(0.016)	

Notes: The table shows the robustness of our estimates of the EIS  $\sigma$  to a number of the assumptions of our structural model. Panel (1) varies the discount factor  $\delta$ . Panel (2) relaxes our assumption that households discount the future with traditional geometric discounting and allows for quasi-hyperbolic  $\beta - \delta$  discounting with present bias parameters  $\beta$  from 0.3 to 1 (geometric discounting). Panel (3) varies the coefficient of relative risk aversion  $\gamma$ . Panel (4) varies (deterministic) future interest rates by shifting the entire yield curve up by 1pp to 3pp. Panel (5) varies real house price growth ranging from house price declines to house price increases an order of magnitude larger. Panel (6) varies the variance of house price growth in household expectations. Panel (7) varies the deterministic lifecycle component of income expectations. In all cases, the lifecycle component is quadratic. The “Peak” parameter gives peak income and the “Slope” parameter gives expected income growth at the time of refinancing. Panels (8) and (9) vary assumptions on income uncertainty, with panel (8) varying the probability of unemployment and panel (9) the unemployment replacement rate.

The message emerging from these robustness checks is that our EIS estimates are insensitive to the parametric assumptions of the model. This confirms the conceptual arguments made earlier that bunching identifies the EIS, because other parameters in general have minor impacts on bunching. It also confirms the simulation exercises in Section 3.3 suggesting that the raw data necessitate a low EIS, because higher values (such as log preferences) would generate a radically different LTV distribution.

Finally, we note that a low EIS could be explained by other behavioural models and frictions than those considered here. For example, [Campbell and Cochrane \(1999\)](#) show that consumption habits may lead to low elasticities of intertemporal substitution. [Chetty and Szeidl \(2016\)](#) show

that consumption commitments (captured by adjustment costs in consumption) may generate similar behaviours as consumption habits. Such models offer potential mechanisms that could explain the low values of the EIS we find.<sup>44</sup>

## 6. CONCLUSION

A large literature estimates the EIS using non-experimental data and structural models. 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. We contribute to efforts to estimate the EIS with a new approach that combines quasi-experimental identification with structural methods.

Our new methodology translates empirical moments arising from bunching at mortgage interest notches into EIS estimates. Using administrative mortgage data from the U.K., we first illustrate our approach in a simple two-period model, and then generalize to a rich stochastic lifecycle model. The two models produce very similar results: the EIS is small, around 0.1. Although we observe lots of bunching at notches, the intertemporal incentives created by those notches are so large that there would have to be much more bunching to justify large values of the EIS. Our results are close to the level obtained in the early macro-based literature (e.g. Hall, 1988), although we have arrived at this conclusion using a fundamentally different approach.

*Acknowledgments.* We thank Orazio Attanasio, Adrien Auclert, Chris Carroll, Raj Chetty, Jeff Clemens, Gordon Dahl, Mariacristina De Nardi, Rebecca Diamond, Roger Gordon, Bob Hall, Ethan Kaplan, Greg Kaplan, Patrick Kehoe, David Laibson, Emi Nakamura, Petra Persson, Ricardo Reis, Jose-Victor Rios-Rull, Emmanuel Saez, Orié Shelef, Jón Steinsson, Gregory Thwaites, Gianluca Violante, Garry Young, Stephen Zeldes, and four anonymous referees for helpful comments and discussion. All charts and estimates use data provided to the Bank of England by the Financial Conduct Authority and MoneyFacts. This research was carried out as part of the Bank of England's One Bank Research Agenda and an earlier draft was circulated under the title "Interest rates, debt and intertemporal allocation: evidence from notched mortgage contracts in the United Kingdom". 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.

### Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

## REFERENCES

- ATTANASIO, O. P. and WEBER, G. (1993), "Consumption Growth, the Interest Rate and Aggregation", *Review of Economic Studies*, **60**, 631–649.
- ATTANASIO, O. P. and WEBER, G. 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, **103**, 1121–1157.
- ATTANASIO, O. P. and WEBER, G. 2010. "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy", *Journal of Economic Literature*, **48**, 693–751.
- BEST, M., CLOYNE, J., ILZETZKI, E., ET AL. 2018. "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches" (Working Paper No. 24948, NBER).
- CAMPBELL, J. and COCHRANE, J. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy*, **107**, 205–251.
- CAMPBELL, J. Y. and MANKIWI, N. G. 1989. "The Response of Consumption to Income: A Cross-Country Investigation", *European Economic Review*, **35**, 723–756.

44. However, note that non-responses due to consumption commitments may be partly captured by the Kleven–Waseem friction adjustment approach.

- CHETTY, R., FRIEDMAN, J. N., OLSEN, T., ET AL. 2010. "Adjustment Costs, Firm Responses, and Labor Supply Elasticities: Evidence from Danish Tax Records" (Working Paper 15617, NBER).
- CHETTY, R., FRIEDMAN, J. N., OLSEN, T., ET AL. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records", *Quarterly Journal of Economics*, **126**, 749–804.
- CHETTY, R. and SZEIDL, A. 2016. "Consumption Commitments and Habit Formation", *Econometrica*, **84**, 855–890.
- DEFUSCO, A. and PACIOREK, A. 2017. "The Interest Rate Elasticity of Mortgage Demand: Evidence From Bunching at the Conforming Loan Limit", *American Economic Journal: Economic Policy*, **9**, 210–240.
- DUNSKY, R. M. and FOLLAIN, J. R. 2000. "Tax-Induced Portfolio Reshuffling: The Case of the Mortgage Interest Deduction", *Real Estate Economics*, **28**, 683–718.
- EINAV, L., FINKELSTEIN, A., and SCHRIMPF, P. 2015. "The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D", *Quarterly Journal of Economics*, **130**, 841–899.
- EINAV, L., FINKELSTEIN, A., and SCHRIMPF, P. 2017. "Bunching at the Kink: Implications for Spending Responses to Health Insurance Contracts", *Journal of Public Economics*, **146**, 27–40.
- EPSTEIN, L. G. and ZIN, S. E. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica*, **57**, 937–969.
- EVANS, O. J. 1983. "Tax Policy, the Interest Elasticity of Saving, and Capital Accumulation: Numerical Analysis of Theoretical Models", *American Economic Review*, **73**, 398–410.
- FOLLAIN, J. R. and DUNSKY, R. M. 1997. "The Demand for Mortgage Debt and the Income Tax", *Journal of Housing Research*, **8**, 155–199.
- GELBER, A. M., JONES, D., and SACKS, D. W. 2019. "Estimating Earnings Adjustment Frictions: Method and Evidence from the Social Security Earnings Test", *American Economic Journal: Applied Economics*. forthcoming.
- GRUBER, J. 2013. "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution", *Quarterly Journal of Finance*, **3**, 1–20.
- HALL, R. E. 1988. "Intertemporal Substitution in Consumption", *Journal of Political Economy*, **96**, 339–357.
- HAVRÁNEK, T. 2015. "Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting", *Journal of the European Economic Association*, **13**, 1180–1204.
- HECKMAN, J. 1979. "Sample Selection Bias as a Specification Error", *Econometrica*, **47**, 153–161.
- JAPPELLI, T. and PISTAFERRI, L. 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.
- KLEVEN, H. 2016. "Bunching", *Annual Review of Economics*, **8**, 435–464.
- KLEVEN, H. J. and WASEEM, M. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan", *Quarterly Journal of Economics*, **128**, 669–723.
- LING, D. C. and MCGILL, G. A. 1998. "Evidence on the Demand for Mortgage Debt by Owner-Occupants", *Journal of Urban Economics*, **44**, 391–414.
- MARTINS, N. C. and VILLANUEVA, E. 2006. "The Impact of Mortgage Interest Rate Subsidies on Household Borrowing", *Journal of Public Economics*, **90**, 1601–1623.
- PIAZZESI, M. and SCHNEIDER, M. 2016. "Housing and Macroeconomics" (Working Papers 22354, NBER).
- SAEZ, E. 1999. "Do Taxpayers Bunch at Kink Points?" (Working Paper 7366, NBER).
- SAEZ, E. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, **2**, 180–212.
- SUMMERS, L. H. 1981. "Capital Taxation and Accumulation in a Life Cycle Growth Model", *American Economic Review*, **71**, 533–544.
- VISSING-JØRGENSEN, A. 2002. "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution", *Journal of Political Economy*, **110**, 825–853.
- WEIL, P. 1990. "Unexpected Utility in Macroeconomics", *Quarterly Journal of Economics*, **105**, 29–42.
- ZELDES, S. P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation", *Journal of Political Economy*, **97**, 305–346.