Erasmus University Rotterdam - Erasmus School of Economics

Master Thesis Policy Economics

Estimating the Elasticity of Mortgage Demand: Mortgage Notches in the Netherlands*

Author: A.J.M. van der Bijl†

Thesis supervisor: Dr. A.A.F. Gerritsen

Second assessor: Dr. S.V. Kapoor

Final version handed in on May 1, 2020.

Abstract

This paper exploits a quasi-experiment in the Dutch mortgage market – households bunch at certain loan-to-value ratio thresholds due to discrete jumps ('notches') in the mortgage interest rate schedule – to estimate the relation between the interest rate and mortgage demand. This is done using a novel empirical methodology called bunching analysis. The estimates indicate that an increase in the mortgage interest rate of 1 percentage point reduces average mortgage demand by 3.5 percent. Analysing specific banks yields semi-elasticities between 1.1 and 7.5. These figures might be valuable when assessing the effect that monetary policies and reductions in the mortgage interest deduction rate (*hypotheekrenteaftrek*) have on the size and composition of mortgage debt in the Netherlands.

*Most of the work on this thesis was done during an internship at the Dutch Ministry of Internal Affairs. As such, this paper could not have been realised had it not been for the invaluable help and guidance provided by Jeske de Lint (BZK), Jurgen van Heemst (BZK) and Mauro Mastrogiacomo (DNB).

†Student number: 415265

Contents

1 Introduction

Households that want to buy a home encounter many tough choices. For instance, they have to strike a balance between financing their dwelling with a mortgage and paying in cash. This financing decision directly affects household finances and the financial well-being of all home-owners. The choices that households make in this respect relate directly to the structure of the financial economy and have implications for macro-prudential policy. For example, if home-owners lend a lot and invest little capital, 'underwater mortgages' and credit risks for banks loom large when the economic tide turns and house prices drop. In spite of this, there is not much research that credibly investigates how home-owners react to changes in the choice parameters that are relevant to this financing decision. Thus, it is not so clear how households alter their choices when interest rates change. This relation between mortgage size and interest rates is important for policy makers to take into account. For example, monetary policy affects the size and composition of mortgage debt through interest rates (Jordà, Schularick, & Taylor, 2015). Therefore, it is beneficial to have credible estimates of how interest rates affect the size and distribution of mortgage debt. What is more, the relation between interest rates and mortgage demand is closely related to the debate on the mortgage interest deduction (MID) – *hypotheekrenteaftrek* in Dutch. MID policies directly affect the effective mortgage interest rate that homeowners face (Poterba, 1984; Poterba & Sinai, 2008). When alterations in the MID rate are contemplated, it is valuable to have causal estimates of how the corresponding change in the effective tax rate influences the mortgage market.[1]

In this paper I focus on the intensive margin response of households to mortgage interest rate fluctuations. By estimating the elasticity of mortgage demand with respect to the mortgage interest rate, I identify to what degree home-owners substitute cash for mortgage debt when confronted with lower interest rates, and vice versa.

Estimating this elasticity is not straight-forward. The reason for this is that endogeneity issues arise when regressing mortgage debt on interest rates. To circumvent this problem, I exploit a quasi-experiment that is present in the Dutch mortgage market. The setting can be sketched as follows. In the Netherlands, the mortgage interest rate is a function of the loan-to-value ratio (LTV). That is, the mortgage interest one has to pay depends on the mortgage amount relative to a dwelling's value: $\text{LTV} = \frac{\text{loan amount}}{\text{house value}}$. An appealing quality of the Dutch interest rate schedule is that it features discrete jumps at some LTV thresholds. An example from a Dutch bank might clarify this. At Rabobank, the mortgage interest rate increases by 0.2 percentage points when a households' LTV surpasses 90%, *irrespective of the other characteristics of the mortgage*. Clearly, this incentivises households to 'bunch' below these LTV thresholds. After all, a slightly higher LTV increases the interest rate on the *entire* mortgage. To illustrate this, consider a dwelling financed by Rabobank that is

[1] In the Netherlands, the tax treatment of owner-occupied housing is a hotly debated topic; many economists have argued for a reduction of the MID (see for instance De Mooij, Van Ewijk, & Jacobs, 2006; Van Ewijk, Jacobs, & De Mooij, 2007; Van Ewijk & Lejour, 2019). Currently, the maximum rate at which mortgages can be deducted is being reduced (Hoekstra, 2019).

worth 300 thousand Euros (roughly the average price of a house). This household could lower its LTV from 90.1% to 90.0% by investing €300 in cash. Doing so, decreases *yearly* interest expenses by roughly €550.[2] This trade-off is favourable for all rational households with realistic intertemporal preferences. Thus, the incentive to bunch is strong. Indeed, Figure 1 shows that households do bunch at the LTV thresholds (see also Figure B.1 in the Appendix).

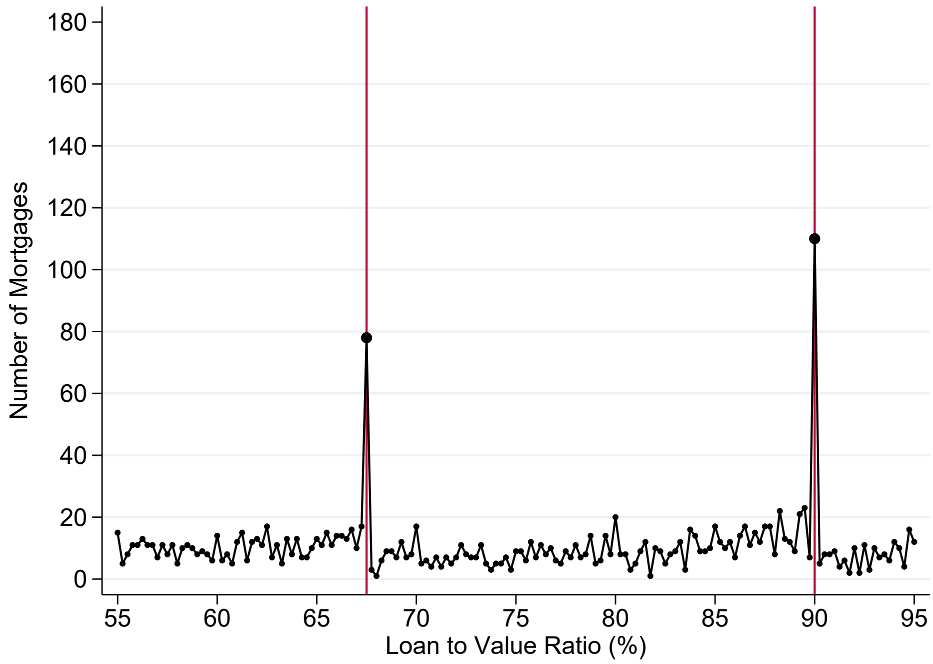

Figure 1: The LTV distribution for Rabobank

Notes: This graph displays the observed distribution of the LTVs among Dutch households that have taken out a mortgage at Rabobank (or an affiliated label) between 2013 and 2018. For this particular bank, interest rate notches are located at the LTV thresholds 67.5 and 90.

In this paper, I use a recent methodological advancement in the public finance literature to translate the peaks from these figures into reduced-form elasticities. This method, commonly referred to as 'bunching analysis', was pioneered by Saez (2010); Chetty, Friedman, Olsen, and Pistaferri (2011); Kleven and Waseem (2013), who focused on taxpayers' behavioural responses to changes in the tax rate.

First, I lay out a theoretical model in which households can choose between consumption and investing in housing. Agents can choose to finance their housing investment with a mortgage, and maximise their utility by picking an LTV. When doing so, households face a 'notched' budget constraint, due to a discrete change in the interest rate at a threshold level of the LTV. It turns out that there exists a marginal buncher who is indifferent between locating at the threshold LTV level and at some point higher up in the LTV distribution. The aim is to retrieve the LTV that this marginal buncher would have taken in the absence

[2] In this calculation I use the fact that the mean interest rate is 2.36% in a small bandwidth just below the 90% notch for Rabobank customers in my sample.

of a notched interest rate schedule.

The difference between this ‘original’ LTV level and the threshold is the behavioural response of the marginal buncher. This response can be used to calculate the reduced-form elasticity of mortgage demand. Because of the fact that the budget constraint is notched, it would be incorrect to retrieve the elasticity by dividing the percentage change in the LTV (i.e., the behavioural response) by the percentage change in the interest rate. Therefore, I present some formulae that allow for an approximation of the elasticity by means of a reduced-form approach.

What is more, I elaborate on the functioning of the bunching estimator. Most important here is that I explain the basic intuition that allows one to solve for the behavioural response of the marginal buncher. Moreover, I show that this behavioural response can be seen as the effect for the *average* marginal buncher. Also, I discuss the potential problem of optimisation frictions and present the most common method in the literature to correct for this.

The Dutch mortgage market has some specific features that makes it a good match for bunching analysis. Not only do banks have mortgage interest rate schedules that feature discrete jumps, but also consumers are fully informed of these details due to recently implemented (in 2013) legislation. As a result, the mortgage landscape in the Netherlands features a quasi-experiment that strongly incentives households to make a specific LTV choice: either at the threshold or somewhere above it.

The data that I use come from the Loan Level Data (LLD) of De Nederlandsche Bank (DNB), the Dutch central bank. Because I want to exploit the Single Interest Rate policy that was implemented in 2013, I only use newly created mortgages that originated after the policy’s introduction. As a result, the dataset becomes dramatically smaller: only 15,000 of the 100,000 mortgages remain to be analysed. As these are distributed over a number of banks who each set different LTV thresholds, I choose to analyse only the biggest banks (Rabobank, ABN AMRO and SNS). This selection slims down the available data to approximately 10,000 unique observations.

To be able to back out the behavioural response of the marginal buncher, it is of paramount importance to have a reliable counterfactual distribution of LTVs. This counterfactual would be the distribution had there not been notches in the mortgage interest rate schedule. I follow the standard approach in the literature, by fitting a ‘flexible polynomial’ through the observed distribution, leaving out an interval around the notch.

Finally, this yields semi-elasticities that range range from -0.011 to -0.075, with a pooled estimate of -0.035. These semi-elasticities should be interpreted as the change in the mortgage balance in percentages given a one percentage point increase in the interest rate. In other words, my estimates of the (absolute) semi-elasticity of mortgage demand are between 1.1% and 7.5%, with a pooled estimate of 3.5%. These estimates are slightly higher than findings in the existing literature. However, it should be noted that this literature is limited in size.

The paper is organised as follows. Section 2 provides a brief introduction to bunching.

It describes how bunching was applied initially, and provides a (non-exhaustive) overview of bunching applications in the literature. Section 3 explicates a theoretical model that characterises the intensive margin financing decision of households. What is more, this section chronicles the theoretical foundations of the bunching estimator in application to the theoretical model. Next, it construes a formula that allows me to calculate the reduced-form elasticity with measurable sufficient statistics. Lastly, it address optimisation frictions and ways to control for them. Section 4 sketches the institutional context of the Dutch mortgage market and describes the data. It also specifies the empirical methodology of this paper. Section 5 reports my estimates and relates them to the literature. The last section concludes.

2 Earlier literature

2.1 A brief introduction to bunching

A recent development in empirical economics is an approach called 'bunching'.[3] This new methodology exploits individuals' bunching behaviour; that is, people's conjoint self-selection into a specific location of some variable's range. When bunching behaviour is present, a histogram plot typically shows seemingly random spikes in the distribution of the relevant variable. Commonly, this explicit, *en masse* positioning is due to discontinuities in the conditions that constrain agent's objective functions. The bunching technique gauges the responses to these discontinuities and uses them to estimate (reduced-form) elasticities or other structural parameters.

There are two distinct conceptual frameworks in the bunching literature. The original approach – the one developed by Saez (2010) and Chetty et al. (2011) – makes use of 'kink points'. A kink point occurs when a budget set has a sudden change in its slope, such that the budget set is still continuous, but not differentiable at the kink point. For instance, if there is a discrete increase in the marginal income tax rate at income level z^*, individuals have an incentive to bunch (i.e., pick an income level z) just below z^*.[4] In this context, bunching analysis can be used to retrieve the compensated elasticity of taxable income by comparing the 'kinked' distribution to an imputed counterfactual distribution (i.e., the distribution as it would look like had there not been a kink in the budget set).[5] Conversely, in the case of a decrease in the marginal tax function, there will be no excess bunching below the corresponding point in the earnings distribution, but a gap in the density function.

The second, more recent bunching framework – pioneered by Kleven and Waseem (2013) – exploits 'notched' budget sets. A notch is present when the budget set features discrete jumps, such that the budget set is discontinuous (and thus non-differentiable) at the notch

[3] See for an accessible introduction to bunching this page by Itai Trilnick (`https://bit.ly/3a5th4c`). For a more technical explanation, see the comprehensive review by Kleven (2016).

[4] In theory, this is the case in any income tax scheme that features income brackets.

[5] For a formal explanation of bunching in the presence of kinks, see Saez (2010) and Kleven (2016) who use a standard (linear) income tax model – similar to the ones often used in the optimal taxation literature – that features workers who are heterogeneous in ability to show what I described concisely.

point. In the context of income taxation a notch corresponds to a discrete increase in the *average* tax function. Note that an increase in the marginal tax function (i.e., kink) at income level z^* means that the higher tax rate applies to every monetary unit earned in excess of z^*. In contrast, an increase in the average tax function (i.e., notch) entails that once earnings exceed z^* the higher tax rate applies to *all* earnings. This implies that a notch induces stronger behavioural responses than a kink, since the former 'creates a region of strictly dominated choice' (Kleven, 2016). Given a jump in the average tax function at z^* there is a 'dominated' space $(z^*, z^* + \Delta z^D)$ in which it is possible to work *less* whilst having *higher* after-tax income by choosing an earnings level $z < z^*$.

Note that bunching is loosely related to other recent empirical frameworks such as regression discontinuity design (RDD) and regression kink design (RKD) (for a detailed description of these techniques, see Imbens & Lemieux, 2008; Card, Lee, Pei, & Weber, 2015). RDD and RKD analysis is similar to bunching in the sense that all three techniques exploit kinked and notched incentives. Notwithstanding this similarity, they differ with respect to the exogeneity of the assignment variable. The identifying assumption in the RDD and RKD methodology is that subjects do not have (perfect) control over the assignment variable. That is, people cannot (fully) manipulate on which side of a given threshold they end up. Contrarily, bunching explicitly requires that individuals can (perfectly) influence the running variable.

The sudden emergence of bunching in the literature coincides with the increasingly prominent use of administrative data. Since bunching responses are typically very local, it is of paramount importance to use large data sets that are free of measurement error, which explains the increasing popularity of bunching designs in the light of a wider availability of administrative data (Kleven, 2016).

2.2 Seminal bunching papers

The pioneers of the bunching technique made their contributions in the fields of labour and public economics, using it to estimate workers' reactions to incentives in the income tax schedule. The first application of bunching can be found in a paper by Saez (2010), who studies individuals who face kink points in the income tax schedule that which are caused by the Earned Income Tax Credit policy (EITC). Tax payers respond to these kinks by bunching just below them. Subsequently, the bunching mass is used to estimate the compensated elasticity of reported income with respect to the marginal income tax rate.

Another seminal paper bunching paper is the one by Chetty et al. (2011). These authors study Danish tax data in which bunching behaviour of taxpayers is present at kink points in the income tax schedule. Translating this behaviour into the elasticity of taxable income with respect to the tax rate, the researchers find that the effects are smaller than expected. That is, one would expect stronger bunching at the kink due to the fact that doing so increases workers' utility. They show that the observed, limited response is most likely the result of frictions in the labour market. Labour supply is for a part shaped by the adjustment costs and hours constraints that workers face. In turn, this means that microeconometric estimates of, for instance, the elasticity of taxable income only account for a part of the 'true' structural elasticity.

In contrast to the first two papers, Kleven and Waseem (2013) study bunching behaviour elicited by notches, instead of kinks. The main contribution of this paper is methodological, as it introduces a framework that allows for the identification of structural and reduced-form elasticities for agents with a notched budget set. They use their framework to analyse behavioural responses of Pakistani self employed workers who face notches in their income tax schedule (i.e., sudden increases of the *average* tax rate). They find that these tax notches elicit sizeable behavioural responses, despite the fact that the structural elasticity that drives these responses is moderate. This finding empirically substantiates that notches are highly distortionary and inefficient in the context of income taxation.

More recently, bunching has been applied to several other fields as well. Table 1 presents a non-exhaustive list of these applications – many of which are published in top-tier journals.

2.3 Directly related literature

In this section I provide a more detailed overview of the papers that are of direct relevance to my work. These studies can be divided in two categories: those that are methodologically related to my paper, and those that provide estimates of the elasticity of mortgage demand.

My survey is directly related to the 'founding' bunching literature. In addition, there is a direct link to some other bunching papers, of which three stand out. With respect to the theoretical background of my research, the work by DeFusco and Paciorek (2017) is most related. The authors use a regulation that is imposed on government-sponsored enterprises. As a result of this regulation, there is a notch in the interest rates of loans at the conforming loan limit of Fannie Mae and Freddie Mac. They use this quasi-experimental variation to estimate the elasticity of mortgage demand. In their paper, the researchers present a theoretical model that is similar to the one in this paper. However, the empirical methodology of their paper and mine is different.[6]

Regarding the empirical methodology, my inquiry benefited greatly from the study by Best and Kleven (2018), which incorporates all of the new developments in bunching, and provides a norm-setting framework for bunching with notches. As such, it is a key reference in the recent public finance bunching literature. The paper exploits quasi-experimental variation due to notched transaction taxes in the British housing market. Using bunching analysis, the authors show that transaction taxes have an impeding effect housing market activity. Stamp duties distort the pricing, quantity and timing of transactions on the housing market.

A recent paper by Best, Cloyne, Ilzetzki, and Kleven (2015, 2020) analyses the UK mortgage market, and is very much related to my work with respect to its institutional context. These researchers model and estimate the elasticity of intertemporal substitution (EIS) using intertemporal substitution preferences of households that they retrieved via

[6]That is, I follow Best and Kleven (2018).

Table 1: A (non-exhaustive) list of papers featuring bunching analysis

Study	Source of bunching	Finding
Kopczuk and Munroe (2015)	Notched transfer tax schedule; dwellings worth more than 1 million USD face a 1% 'mansion tax' upon sale	Transfer taxes distort the housing market by affecting prices and sales
Brown (2013); Manoli and Weber (2016)	Kinks in the retirement benefit scheme; notched severance payment schedule	Estimates of the effect of retirement benefits on workers' extensive margin labour market decisions
Brehm et al. (2017)	Monetary incentives at thresholds in a scale of measurable teacher productivity	The performance pay schedule does not successfully elicit workers' productivity growth
Cengiz et al. (2019)	Increase of the minimum wage causes a spike of jobs at the minimum wage in the wage distribution	The effect of minimum wages on low-wage jobs; no jobs are lost since missing mass equals bunching mass
Søgaard (2019)	Notch in the income tax rate schedule for Danish students	Optimisation frictions in workers' labour supply prevent bunching behaviour
Finkelstein et al. (2019)	Notches in a US healthcare insurance subsidy scheme	An estimation of the willingness to pay for insurance
Ruh and Staubli (2019)	Notches in the Austrian disability insurance scheme	The effect of financial incentives on the earnings of disability insurance beneficiaries
Einav et al. (2015, 2019)	Kink in the 'out of pocket' costs of a health insurance scheme	The response of drug demand to incentives in pricing schedules
Best and Kleven (2018)	Notches in the 'stamp duty' schedule (land taxes) for houses in the UK	The effect of a transfer tax on price, volume, timing and other housing market characteristics
Diamond and Persson (2016)	High school grades spike at certain levels due to teacher indiscretion	Earnings development of high school students given bunching variation in high school grades
Keys and Wang (2019)	Kink in the credit card minimum payment schedule	The relation between minimum credit card payments and liquidity constraints
Goupille-Lebret and Infante (2018)	Notches in the French inheritance tax rate	The relation between inheritance decisions and the inheritance tax rate
Jakobsen et al. (2018)	Kink in the Danish wealth tax scheme	The effect of wealth taxation on the accumulation of wealth
Sallee and Slemrod (2012); Ito and Sallee (2018)	Notched fuel tax schedule	Car producers' supply decisions in response to fuel policies
Almunia and Lopez-Rodriguez (2018)	Notch in Spanish firms' revenue tax schedule	The response of reported firm revenue to governmental tax compliance policies
Bosch et al. (2019)	Kinks and notches in household budget sets due to tax transfers	Absence of behavioural responses to kinks and notches in the Dutch allowance system

bunching. As these authors exploit notches in the mortgage interest rate schedule, there is an obvious connection to my work. However, they explicitly only assess households that refinance their mortgage – due to a peculiarity in the UK mortgage market, most households refinance every 3-5 years. Moreover, they do not estimate the elasticity of mortgage demand, but use the variation in interest rates to update a structural model of household choice to retrieve the EIS.

With respect to the relation between interest rates and mortgage size that I estimate, a small literature exists. The reason for the limited size of this literature is that there exists basically no exogenous variation in mortgage interest rates. Since both interest rates and house prices depend on confounding determinants such as the state of the economy and dwelling characteristics, simply regressing mortgage size or loan-to-value ratio (LTV) levels on interest rates yields estimates that suffer from omitted variable bias.

First of all, Fuster and Zafar (2015) analyse the results of a 'strategic survey' on the sensitivity of housing demand to several parameters. In their research, homeowners answer questions on their willingness to pay for a dwelling similar to their own under – among other things – different mortgage interest rates. Their findings imply an interest rate semi-elasticity of mortgage demand between 0.6% and 1.8%. Moreover, the earlier mentioned paper by DeFusco and Paciorek (2017) reports semi-elasticities of 1.5% to 2%.

Furthermore, there exists a modest literature that attempts to use variation in after-tax mortgage interest rates caused by government policies (e.g., mortgage interest deduction (MID), subsidies, regulatory policies etc.) to calculate mortgage demand elasticities. For example, Follain and Dunsky (1997) examine the effect of a reduction in the mortgage deduction rate on mortgage size. Since they report the elasticity of mortgage debt with respect to the MID rate, the estimate cannot be compared directly to my estimates of the (semi-)elasticity. Martins and Villanueva (2006) investigate a reform of a government program that subsidises mortgage expenses for lower income households. The problem with this paper is that the authors do not measure the elasticity of mortgage demand on the intensive margin, but the probability that homeowners do borrow given an interest rate change. Another example in this strain of literature is by Jappelli and Pistaferri (2007) who evaluate a reduction in the MID rate in Italy. Surprisingly – and in contrast with theoretical predictions as well as the empirical literature – they do not find that the higher cost of mortgage borrowing affects the demand for mortgage debt on the extensive or intensive margin. A possible explanation for this is a lack of financial information and credit rationing during the reform period (1992-1994).

In the non-housing literature, Attanasio, Goldberg, and Kyriazidou (2008) find semi-elasticities for automobile financing schemes. Karlan and Zinman (2008), who look at consumer micro-credit, also report interest rate elasticities of credit demand. Both findings fall in the same range as the semi-elasticities found by Fuster and Zafar (2015) and DeFusco and Paciorek (2017).

3 Theoretical framework

3.1 A model of LTV choice

In the spirit of the models by Brueckner (1994) and DeFusco and Paciorek (2017) I present a two-period model in which households make an intensive margin choice to finance their housing with a mortgage. The mechanism of my specification is identical to the one by DeFusco and Paciorek (2017). The only difference between our models is that their choice variable is the absolute mortgage size, whereas my choice variable is the mortgage size relative to the value of the dwelling – since I am interested in the loan-to-value ratio (LTV). This means that normalising house values in my specification (i.e., setting the house value equal to one) yields the original model.

The type of model presented in this paper is quite standard in the bunching literature. Hence, it bears resemblance to the theoretical models used in the seminal bunching papers, such as those by Saez (2010), Chetty et al. (2011), Kleven and Waseem (2013) and Best et al. (2020). Despite their stylised nature, these models provide some helpful insights into the mechanisms at play in a bunching setting.

Consider a world in which households live for two periods. In the first period, households receive (labour) income, which they subsequently invest in owner-occupied housing or spend on non-housing consumption goods. Households have the option to finance housing investment with a mortgage. In the second period households liquidate their homes, return the mortgage principal, transfer the mortgage interest payable and consume the remainder of their wealth. The first period in this world can be characterised as 'working life' and the second period as 'retirement'. For simplicity, there exists no financial sector other than the mortgage market. That is, all saving takes place through the financing decision that underlies owner-occupied housing. Here I implicitly assume that there is no intertemporal substitution from 'retirement' to 'working life' (i.e., households do not 'borrow' from the second period through negative mortgages). In addition, I make the simplifying assumption that dwellings do not generate capital gains or losses (i.e., the value of the house does not change throughout the two periods). These assumptions do not alter the mechanism that I study with the model of LTV choice. The theoretical specification provides an overview of the relevant choice parameters, and will serve as the theoretical basis for the bunching estimator.

An LTV ratio is the fraction of a dwelling's value that is financed with a mortgage. Typically, households acquire a dwelling with market value p, which they finance with a mortgage sized m. Hence, the LTV is given by the fraction $\ell = \frac{m}{p}$. Note that the purchase price of the dwelling and its market value are not necessarily the same. The purchase price is simply the price that the seller and buyer have agreed on. Conversely, the market value is established by a professional appraiser (*taxateur*). This value is based on an array of characteristics of the dwelling (e.g., location and general condition of the house, it's size, etc.), and is thus not directly influenced by the price negotiation. For this reason, the market value is fixed: it cannot be altered by the household, since it is objectively 'given'.

This fixed market value is the denominator of the LTV ratio. Thus, upon acquiring a dwelling, households can only influence ℓ through m. To put this differently: the only way for households to lower their LTV with the aim of ending up in a lower mortgage interest bracket, is taking out a smaller mortgage. There is, however, one problem that looms here. It is possible that appraisers – who know about the optimisation problem that the households face – try to 'help out'. In theory, they could overestimate a dwelling's market value in order to deflate the LTV ratio. There is no concrete evidence whether this actually occurs (nor how often), but the theoretical possibility is a caveat that has to be taken into consideration.

Households derive utility from the consumption of a non-housing good c. All agents maximise their lifetime utility by selecting non-housing consumption in both periods, so that their objective function is specified as

$$\max_{c_1, c_2} \{\mathcal{U}(c_1, c_2) \equiv u(c_1) + \delta u(c_2)\}, \tag{1}$$

where $\delta \in (0, 1)$ is the rate at which households discount future utility. In accordance with the bunching literature, I make the explicit assumption that household preferences behave according to the quasi-linear, isoelastic utility function

$$u(c) \equiv \frac{1}{1-\sigma} c^{1-\sigma}. \tag{2}$$

Every household enters the first period with a wealth endowment and earns some labour income, the sum of which I call W_0. As a result, after the housing investment decision, consumption in period one can be written as

$$c_1 = W_0 - (1 - \ell)p. \tag{3}$$

In the second period housing is liquidated and the mortgage is paid off, which constrains period-two consumption to

$$c_2 = p[1 - \ell(1 + r)], \tag{4}$$

where r is the mortgage interest rate.

To be able to use the model in a bunching context, some additional assumptions are necessary. I follow DeFusco and Paciorek (2017), who introduce heterogeneity in preferences by allowing agents to have different discount rates. To ensure a smooth distribution of household preferences, I make the additional assumption that the discount factors are assigned in such a way that their density function $f(\delta)$ is continuous and differentiable. Moreover, I introduce the assumption that households are homogeneous in period-one income W_0, house value p and elasticity σ. Assuming homogeneity in these parameters allows me to solve for the unique LTV, ℓ^*, that maximises the total utility of a household. For this derivation, see section A.1 in the appendix. The LTV level that maximises total utility is given by

$$\ell^* = \frac{1 + (1 - \frac{W_0}{p})\,(\delta(1+r))^{\frac{1}{\sigma}}}{(1+r) + (\delta(1+r))^{\frac{1}{\sigma}}}. \tag{5}$$

Since I assumed that W_0, p and σ are the same for all households, equation (5) posits a one-to-one relationship between the (heterogeneous) discount rates of households and their optimal LTV ratios. In the absence of mortgage interest rate notches this relationship brings about an LTV density function which I will refer to as $f_0(\ell)$.

Allowing for heterogeneity in W_0, p and σ does not alter any of the mechanisms of the model but makes it impossible to solve for ℓ^* algebraically. Nevertheless, it is possible to estimate the bunching effect whilst allowing for heterogeneity in these parameters – I show how this alters the bunching estimation in the section that specifies the bunching estimator (3.4).

3.2 Mortgage interest rate as a step-function

To proceed, suppose that the mortgage interest rate schedule has a notch (i.e., increases discontinuously from r to $r + \Delta r$) at the LTV threshold ℓ^τ. It follows that the interest rate is now no longer constant, but a step-function of the LTV, such that

$$r(\ell) = r + \Delta r \times \mathbf{I}\left[\ell > \ell^\tau\right], \tag{6}$$

where $\mathbf{I}[\cdot]$ is an indicator function. As a result of this notch, the period-two budget constraint from equation (4) becomes

$$c_2^n = p[1 - \ell(1 + r(\ell))] = p[1 - \ell(1 + r + \Delta r \times \mathbf{I}\left[\ell > \ell^\tau\right])]. \tag{7}$$

Total consumption over both periods $\mathcal{C}$ is given by the sum of equations (3) and (7), such that

$$\mathcal{C} = W_0 - \ell p\left(r + \Delta r \times \mathbf{I}\left[\ell > \ell^\tau\right]\right). \tag{8}$$

Note that the budget constraint expressed in equation (8) features a 'proportional notch'. This means that there is a discrete change in the interest liability at the cut-off ℓ^τ (i.e., the budget constraint shifts down by $\ell^\tau \Delta r p$), as well as a discrete change in the interest rate the cut-off (i.e., the slope of the budget constraint decreases by Δr). The term 'proportional notch' was coined by Kleven and Waseem (2013), who distinguish between 'proportional' and 'pure' notches. The latter type would only cause a discrete jump in the interest liability, without affecting the interest rate at the cut-off.

Panel A of Figure 2 shows an household's budget constraint (lifetime consumption $\mathcal{C}$ as a function of LTV level ℓ) as well as two indifference curves that belong to the 'marginal buncher'. Since period-one and period-two consumption are determined by ℓ, these indifference curves correspond to all the possible combinations of c_1 and c_2^n that yield the same utility level. The marginal buncher is the household that is indifferent between an LTV at the threshold ℓ^τ or at the point ℓ^I above the cut-off. Since households are heterogeneous in their discount rates, the marginal buncher is the household with the smallest δ in the group of bunching households.[7] In the absence of the interest rate notch, this household would have positioned itself at $\ell^{\tau} + \Delta\ell$. This means that households that would have had an LTV in the domain $(\ell^{\tau}, \ell^{\tau} + \Delta\ell]$ in a no-notch setting, will bunch at the threshold in the presence of a notch. This implies, of course, that households who would – in a no-notch setting – have an LTV to the right of $\ell^{\tau} + \Delta\ell$ do not bunch. However, they reduce their LTV marginally whilst staying to the right of the threshold. The reason for this is that borrowing money is now more expensive, so that non-bunching households reduce their LTV ('substitution effect').

Panel B of Figure 2 shows the hypothetical densities $f_1(\ell)$ and $f_0(\ell)$. The former is the LTV distribution in the presence of an interest notch, whereas the latter connotes the counterfactual distribution of LTVs if no notch is present. The observed bunching and missing mass in the density function $f_1(\ell)$ are a result of the phenomena described above.

[7] The reason for this is that *of the households that bunch*, the marginal buncher gets the smallest utility reduction. Recall that interest payment takes place in the second period, which households discount with discount factor δ.

Panel A: *The budget set of the marginal bunching household given a notch in the mortgage interest schedule*

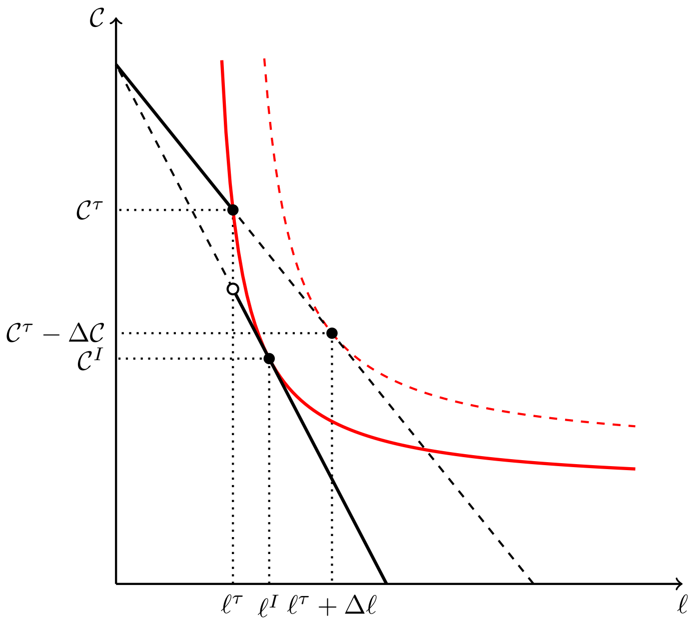

Panel B: *The distribution of the induced 'actual' and counterfactual LTVs*

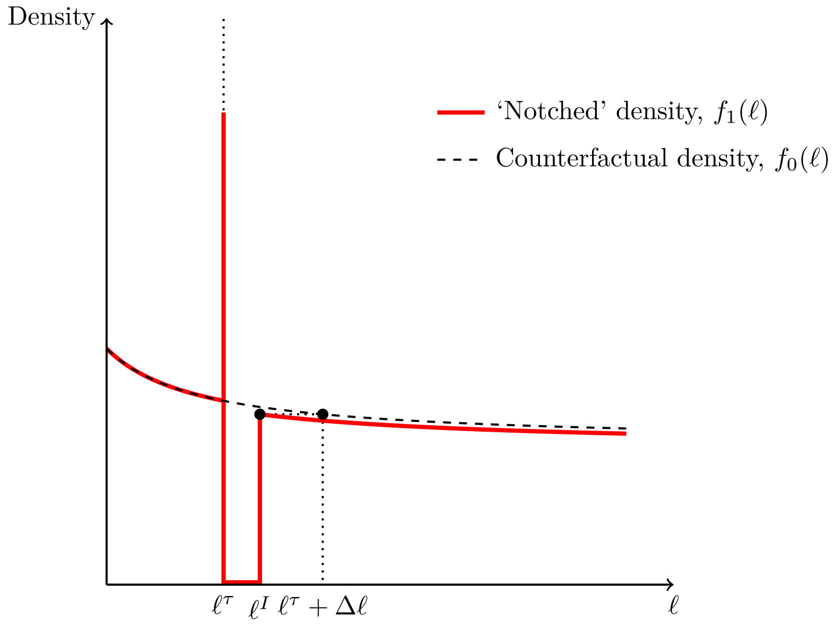

Figure 2: The bunching moment in a model of LTV choice

Notes: The graph in Panel A denotes the budget set of a household that makes a financing choice for an investment in housing. The vertical axis corresponds to lifetime consumption $\mathcal{C}$; the horizontal axis shows LTV levels ℓ. The black line represents the budget constraint – as defined in equation (8) – with slope $-r$ below the threshold ℓ^τ and $-(r + \Delta r)$ above. At the cut-off, the budget constraint is shifted down by $\ell^\tau \Delta rp$. The indifference curves belong to the marginal buncher, the household that is indifferent between staying at a higher LTV (ℓ^I) and bunching below the threshold (ℓ^τ). In the absence of interest rate notches this household has an LTV of $\ell^\tau + \Delta\ell$. Households to the right of $\ell^\tau + \Delta\ell$ reduce their LTV marginally, but stay above the threshold. The interval (ℓ^τ, ℓ^I) is strictly inferior to points at or below ℓ^τ, which means that no household will position itself in this range when a notch exists. The bunching and missing mass caused by this behaviour are depicted in Panel B.

3.3 Reduced-form elasticity of mortgage demand

In the context of a structural model it is possible to back out the elasticity of mortgage demand, provided that there are no unkown parameter values. Recall that the marginal buncher is indifferent between taking an LTV level exactly at the notch (ℓ^τ) and at some point ℓ^I, where $\ell^I \in (\ell^\tau, \ell + \Delta\ell)$. This means that ℓ^τ and ℓ^I both lie on the marginal buncher's preferred indifference curve. As a result, the utility levels that correspond to these two points are equal. This insight yields the indifference equation

$$\mathcal{U}^N - \mathcal{U}^I = 0, \tag{9}$$

which implicitly defines the (compensated) elasticity of mortgage demand σ.[8] See Appendix A.2 for the general derivation and full specification of the equation above. Be that as it may, the difference equation contains the constants W_0, p and δ, of which the values are unknown. This means that it is not possible to retrieve the elasticity of mortgage demand using equation (A.8).

Therefore, it is necessary to develop a non-parametric, reduced-form approach to approximate the elasticity of mortgage demand. I do so by applying the suggestions made in the seminal paper by Kleven and Waseem (2013).[9] The implicit interest calculation in the paper by DeFusco and Paciorek (2017) is similar to mine. The main difference is that these authors are not interested in the LTV level, but in the mortgage size.[10] Moreover, they derive their specification in a different way than I do.[11] Despite these differences, the specifications of Kleven and Waseem (2013), DeFusco and Paciorek (2017) and myself are all based on the same intuition.

The following thought exercise applies this intuition to my theoretical setting. If the marginal buncher picks an LTV at the cut-off, his interest liability equals $R(\ell^\tau) = \ell^\tau rp$. In the hypothetical case that the marginal buncher does not alter his LTV in response to the notch, his interest liability equals $R(\ell^\tau + \Delta\ell) = (\ell^\tau + \Delta\ell)(r + \Delta r)p$. This implies that the marginal buncher, who was originally positioned at LTV level $\ell^\tau + \Delta\ell$, has to pay an interest premium to stay at that same LTV level when confronted with the interest notch. This premium is the difference between $R(\ell^\tau + \Delta\ell)$ and $R(\ell^\tau)$,

$$\text{Interest premium} = R(\ell^\tau + \Delta\ell) - R(\ell^\tau) = \Bigg(\underbrace{\ell^\tau \Delta r}_{a} + \underbrace{\Delta\ell(r + \Delta r)}_{b}\Bigg)p. \tag{10}$$

[8] $\mathcal{U}^N$ and $\mathcal{U}^I$ stand for the total utility at the notch point and the 'indifference' point, respectively.

[9] These authors calculate the increase in income tax that the marginal buncher pays if he does not reduce his earnings z from $z^\tau + \Delta z$ to the threshold, z^τ. They relate this tax increase to the earnings response Δz, such that $t^* = \frac{T(z^\tau + \Delta z) - T(z^\tau)}{\Delta z} \approx t + \frac{\Delta t \cdot z^\tau}{\Delta z^\tau}$ (the approximation is due to the assumption that Δt is small). This is what they call the 'implicit marginal tax rate' – which is equal to what would be the 'implicit average tax rate'. See section *II.B: A reduced-form approximation of the earnings elasticity* (Kleven & Waseem, 2013).

[10] As described in section 3.1, normalising the market value of a dwelling (i.e., setting $p = 1$), yields their specification.

[11] For a comparison, see the section called *Calculating elasticities with a notched budget constraint* (DeFusco & Paciorek, 2017).

Part a of equation (10) corresponds to the 'pure' part of the notch: locating above the cut-off generates an extra interest liability of Δr times the loan size below the cut-off. This is depicted in Panel A of Figure 2, where the budget constraint shifts down by $\ell^\tau \Delta r p$ at the cut-off. Moreover, part b of equation (10) conveys that the notch is not only 'pure', but also 'proportional'. That is, the marginal *and* average interest rate payable on the part of the loan taken out in excess of the cut-off increases by Δr. This can also be seen in Panel A of Figure 2, where the slope of the budget constraint is given by $-(r + \Delta r)$ for LTV levels bigger than ℓ^τ.

This was the first step of the thought exercise. I have established that if the marginal buncher does not respond to the notch, he has to pay the interest premium given by equation (10). The next step is to divide the interest premium by the loan amount in excess of the cut-off, $\Delta\ell p$. This yields the 'implicit' marginal *and* average interest rate that the marginal buncher pays on the '$\Delta\ell p$-part' of the loan. I call this implicit rate $r^\imath$, such that

$$r^\imath \equiv \frac{R(\ell^\tau + \Delta\ell) - R(\ell^\tau)}{\Delta\ell p} = \frac{(\ell^\tau \Delta r + \Delta\ell(r + \Delta r))p}{\Delta\ell p} = r + \Delta r + \frac{\ell^\tau \Delta r}{\Delta\ell}. \tag{11}$$

Equation (11) shows that the implicit marginal interest rate has a constant component equal to the 'post-notch' interest rate $r + \Delta r$, as well as a term $\frac{\ell^\tau \Delta r}{\Delta\ell}$, which decreases in the behavioural response. For infinitesimal behavioural responses, the implicit marginal rate explodes to infinity, $\lim_{\Delta\ell\downarrow 0} r^\imath_p = \infty$. Conversely, for big behavioural responses the implicit marginal rate approaches the post-notch interest rate, $\lim_{\Delta\ell\to\infty} r^\imath_p = r + \Delta r$.

The implicit interest rate $r^\imath$ essentially treats the downward shift of the budget constraint (i.e., the 'pure' part of the notch) as if it were a change in the marginal and average interest rate. To put this differently: the implicit rate regards the behavioural response $\Delta\ell$ as if it were induced by a kink. This kink incorporates the 'pure' component of the notch (part a) in the 'proportional' component of the notch (part b). Adding the implicit interest rate to the budget set illustrates this – see the blue line segment in Figure 3. This figure shows that the 'notched' discontinuity in the budget set is replaced by a 'kinked' budget line with a slope that is equal to the implicit interest rate, $-r^\imath$.

This 'implicit', hypothetical budget line introduces new interior points that are strictly preferred to ℓ^τ, which means that the marginal buncher would not have moved to ℓ^τ under this budget set. Hence, the implicit interest rate (and thereby also the reduced-form elasticities) serve as an upper-bound since they slightly misrepresent (overestimate) the behavioural response. A true functional form of utility would be necessary to retrieve the exact size of the bias – but this would make the reduced-form approximation redundant, since a structural model would mitigate the necessity of the reduced-form approach.

Having established the specification of the implicit marginal rate, I can define the reduced-form elasticity of mortgage demand with respect to the implicit mortgage rate,

$$\varepsilon \equiv \frac{\Delta\ell/\ell^\tau}{\Delta r^\imath / r^\imath} = \frac{\Delta\ell/\ell^\tau}{(r^\imath - r)/r^\imath} = \frac{\Delta\ell}{\ell^\tau}\left(1 + \frac{r}{\Delta r}\frac{\Delta\ell}{\ell^\tau + \Delta\ell}\right). \tag{12}$$

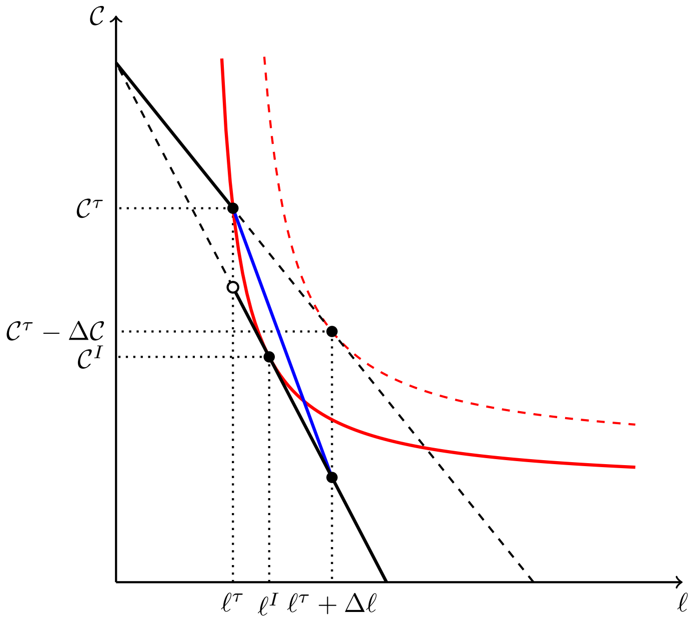

Figure 3: The implicit marginal interest rate in the budget set

Notes: The figure displayed here is similar to Panel A of Figure 2; for an additional explanation see the notes there. In addition to the previous figure, this graph features the hypothetical budget constraint imposed by the implicit marginal interest rate $r^\imath$. The implicit rate is denoted by the blue line, with slope $-r^\imath = -(r + \Delta r + \ell^\tau \Delta r/\Delta\ell)$. This shows that the implicit marginal rate essentially treats the behavioural response to the notch as if it was induced by a kink. Since a part of the blue line is strictly superior to ℓ^τ and ℓ^I (i.e., the marginal buncher can move to a higher indifference curve by locating on an interval on the blue line) the reduced-form approach is clearly an approximation that leads to a bias under some conditions.

Under the assumption that prices are fixed, this equation denotes the elasticity of mortgage demand, as well as 'LTV demand', since $\Delta\ell/\ell^\tau = p\Delta\ell/p\ell^\tau$. Thus, ε shows the percentage change in mortgage size given that the mortgage interest rate changes by one percent. The interpretation of this elasticity is not straightforward; it is more intuitive to display the change of the interest rate in percentage points. Hence, I introduce the semi-elasticity ε_s,

$$\varepsilon_s \equiv \frac{\Delta\ell/\ell^\tau}{\Delta r^\imath} = \frac{\Delta\ell/\ell^\tau}{r^\imath - r} = \frac{\Delta\ell}{\ell^\tau}\frac{\Delta\ell/\Delta r}{\Delta\ell + \ell^\tau}. \tag{13}$$

Lastly, it should be noted that it is incorrect to calculate the elasticity and semi-elasticity by using $\frac{\Delta\ell/\ell^\tau}{\Delta r/r}$ and $\frac{\Delta\ell/\ell^\tau}{\Delta r}$, respectively. These approaches do account for the change in the slope of the budget constraint, but not for the downward shift of the budget constraint. As a result, these expressions lead to estimates that are $\frac{(\ell^\tau+\Delta\ell)r}{(\ell^\tau+\Delta\ell)\Delta r+\Delta\ell r}$ and $\frac{\ell^\tau+\Delta\ell}{\Delta\ell}$ times larger than the results in equation (12) and (13).

3.4 Bunching estimator

When bunching is a phenomenon as displayed in Figure 2, there will be a peak in the density function at ℓ^τ, accompanied by a gap in the distribution on the interval (ℓ^τ, ℓ^I) and a shift to the left of households above $\ell^\tau + \Delta\ell$. In this case the observed bunching behaviour B (i.e., the number of households that bunch) is given by

$$B = \int_{\ell^\tau}^{\ell^\tau + \Delta\ell} f_0(\ell)\ \mathrm{d}\ell \simeq f_0(\ell^\tau)\Delta\ell. \tag{14}$$

Essentially, this means that estimates of observed bunching $\hat{B}$ and the counterfactual distribution $\hat{f}_0(\ell)$ around ℓ^τ suffice to solve for the behavioural response $\Delta\ell$ of the marginal bunching household. Successively, one can use this behavioural response to derive the reduced-form elasticity of mortgage demand with respect to the interest rate.

Recall from the last part of section 3.1 the assumption that the sum of wealth and income W_0, house value p, and the elasticity σ are identical for all households in the model. Allowing for heterogeneity with respect to these variables complicates the analysis slightly. To show that it is still possible to interpret the bunching setting when heterogeneity is introduced, I closely follow Kleven and Waseem (2013) and DeFusco and Paciorek (2017). Consider a joint density function of the aforementioned variables, along with the discount factor that was already assumed to be heterogeneous, $\dot{f}(\delta, \sigma, W_0, p)$. Like δ, the arguments of this function are located on intervals between zero and an upper bound, such that $\sigma \in (0, \bar{\sigma}]$, $W_0 \in (0, \bar{W_0}]$ and $p \in (0, \bar{p}]$. For all unique combinations of the now heterogeneous variables (σ, W_0 and p), one can separately back out the behavioural responses, as in equation (14). This yields the behavioural responses $\Delta\ell_{\sigma,W_0,p}$ of the marginal buncher for all unique combinations of σ, W_0, p. To retrieve the average response of all marginal bunchers, an extra step is necessary.

Consider the counterfactual joint density function $\dot{f}_0(\ell, \sigma, W_0, p)$, which connotes the joint density of LTVs, elasticities, wealth and income, and house values in the absence of an interest rate notch. Next, I define $\tilde{f}_0(\ell)$ as the counterfactual unconditional LTV density, meaning that $\tilde{f}_0(\ell) \equiv \int_\sigma \int_{W_0} \int_p \dot{f}_0(\ell, \sigma, W_0, p)\ \mathrm{d}p\ \mathrm{d}W_0\ \mathrm{d}\sigma$. Under this specification, the bunching behaviour can be denoted as

$$B = \int_\sigma \int_{W_0} \int_p \int_{\ell^\tau}^{\ell^\tau + \Delta\ell_{\sigma,W_0,p}} \dot{f}_0(\ell, \sigma, W_0, p)\ \mathrm{d}\ell\ \mathrm{d}p\ \mathrm{d}W_0\ \mathrm{d}\sigma \simeq \tilde{f}_0(\ell^\tau) E\left[\Delta\ell^\tau_{\sigma,W_0,p}\right]. \tag{15}$$

It is possible to retrieve the *mean* behavioural response $E\left[\Delta\ell^\tau_{\sigma,W_0,p}\right]$ from this formula. This is the average response of all marginal bunchers for all unique combinations of σ, W_0, p.

To be able to back out the behavioural responses $\Delta\ell$ and $E\left[\Delta\ell^\tau_{\sigma,W_0,p}\right]$ from the approximation (right part) of the specifications (14) and (15), it is necessary that $f_0(\ell)$ and $\tilde{f}_0(\ell)$ are constant on the bunching interval $(\ell^\tau, \ell^\tau + \Delta\ell)$. Assuming that the counterfactual density is constant at the threshold allows for a simple application of bunching analysis that does not require much coding.[12]

Be that as it may, this 'linearity of counterfactual' assumption does not necessarily hold in reality, especially when the behavioural response is large. When empirically implementing this approach, it is possible to drop this assumption and use the non-linear or curved

[12]There are some other advantages specific that make this assumption appealing, which I describe in section 4.3).

shape of the counterfactual density. See section 4.3 for more details and a description of my 'non-linearity of counterfactual' approach.[13]

3.5 Optimisation frictions

Thus far, the theoretical framework has been based on the assumption that households will successfully bunch below the LTV threshold whenever this is utility maximising. However, in reality there could be optimisation frictions that prevent households from doing so. For instance, misperception, inattention or (non-mortgage) credit and liquidity constraints could block households from picking an optimal LTV. If these frictions exist, they should be accounted for to prevent a downward bias in the estimate of the behavioural response.

In their seminal paper, Kleven and Waseem (2013) developed a method to adjust for frictions non-parametrically. In the notched income taxation schedule that they study there are dominated regions just above wage notches in which a reduction of gross labour income causes an increase in net labour income. Agents located within these dominated regions are presumed prone to optimisation frictions. Under the assumption that the fraction of individuals impeded by frictions is the same outside the dominated zone, they correct the estimated behavioural responses.

Unfortunately, the household choice that I study does not feature clear-cut dominated regions. Locating above a notch comes at a cost (less future consumption) but brings about a small benefit (a little more current consumption). Whether or not this is optimal depends on a set of parameters.[14] Clearly, for perfectly impatient households ($\delta = 0$) it is optimal *not* to bunch – except in the situation where the threshold LTV maximises period-one utility. For households with $\delta > 0$ this is not necessarily the case. Be that as it may, clearly a region that is dominated for all households does not exist.

Therefore, I implement a more parametric version of the method by Kleven and Waseem (2013), by ruling out extreme preferences and recovering the degree of optimisation frictions from a very small bandwidth above LTV thresholds. The assumption here is that for these households the additional equity injection is so small, and the resulting decrease in interest expenses so high, that their decision must be the result of an optimisation friction. For example, the average value of dwellings around the 67.5 notch (Rabobank customers, 2.5pp bandwidth) is approximately €350,000. For this house price, households with a slightly higher LTV of 67.6 need to invest an additional €350 in order to lower the *yearly* interest expenses by almost €475 ($\Delta r = 0.2$pp). Households failing to inject this equity are presumed prone to optimisation frictions.

Having established this, I divide the observed density mass by the counterfactual density mass in a bin (0.25pp) just above the notch. This gives me the fraction a of non-optimising households in a region that is most likely dominated. Next, I assume that a proxies the

[13] Note that later in this paper I use the terms 'bunching-hole approach' and 'convergence approach' to refer to the 'linearity of counterfactual' and 'non-linearity of counterfactual' methods.

[14] It is only optimal to locate above a notch when $(W_0 - (1 - \ell^{\tau+})\, p)^{1-\sigma} + \delta\, (p\, [1 - \ell^{\tau+} (1 + r + \Delta r)])^{1-\sigma} > (W_0 - (1 - \ell^{\tau})\, p)^{1-\sigma} + \delta\, (p\, [1 - \ell^{\tau} (1 + r)])^{1-\sigma}$, where $\ell^{\tau+} \in (\ell^{\tau}, \bar{\ell})$ is an LTV position above a notch.

(observable) fraction of households prone to frictions further in the distribution, as is common in the literature (see for instance Kleven & Waseem, 2013; Best & Kleven, 2018; Best et al., 2020). In this context, friction-adjusted bunching is given by $\frac{B}{1-a}$ and the adjusted behavioural response by $\frac{\Delta\ell}{1-a}$.

This adjustment serves as an upper bound when controlling for optimisation frictions. Generally, a is not constant but increases in the distance to the notch, since the equity injection necessary grows at a higher rate than the interest gain. However, the advantage of this method is that no parametric assumptions have to be made. Moreover, it allows for an agnostic stance on the sources of the frictions.

4 Empirical framework

4.1 Institutional context

The Dutch mortgage market has some distinctive features that facilitate the bunching approach. Firstly, all banks have notched mortgage interest rate schedules, meaning that interest rates jump discretely at specific levels of the loan-to-value ratio (LTV). When a household exceeds a given LTV threshold, the interest rate on the *entire* mortgage increases. The reason that banks set different interest rates has to do with the fact that mortgages with a higher LTV have a larger propensity to cause losses upon default. Hence, mortgages are categorised in a risk category based on their LTV.[15] Of course, banks want to be compensated for the additional risk that they face when selling a high-LTV mortgage. Therefore, they set a higher mortgage interest rate for mortgages with a high LTV.

The question I have not found a definitive answer to is why banks give their increasing interest schedule a *notched* structure. It would be reasonable to assume that default rates increase smoothly, which would justify a smoothly increasing interest rate schedule. Then, why do banks define the interest rate as a step function of the LTV? The most plausible answer – which is also held by the LTV experts of the Housing Market Department at the Dutch Ministry of Internal Affairs – is that notches merely simplify the process of communicating relevant interest rates to customers.

The second reason why the Dutch mortgage context allows for bunching analysis is that information about the mortgage interest rates is widely available and easy to access. This allows households to make well-informed decisions on their mortgage type and LTV. Moreover, mortgage suppliers cannot offer unique deals to specific households due to government regulation. The relevant legislature in this respect is the Single Interest Rate policy (*Eensporig Rentebeleid*) which has been in effect since 2013. The most important implication of this legislation is set out in article 81a of the Decree on Conduct of Business Supervision of Financial Undertakings (*Besluit Gedragstoezicht financiële ondernemingen*, BGfo), which is a branch of the Dutch Financial Markets Supervision Act (*Wet op het financieel toezicht*,

[15]Banks are obliged to categorise their mortgage portfolio in risk brackets by the Basel III regulation. See for instance this summary of the current regulation by the Basel Committee on Banking Supervision (`https://tinyurl.com/y6wpedle`).

Wft).[16] Article 81a BGfo states that households with a similar risk profile (i.e., LTV level) should be offered equal financing terms.[17] In a publication brief to this law it is explained that this means that new and unique mortgage offers are allowed, under the condition that they are extended (and communicated) to all existing and new customers. As a result, financial institutions that set different mortgage interest rates have to offer the same LTV premiums to all customers. In extension to this, article 51b BGfo requires banks to publish all actual offers and rates on their website. See Figure B.2 in the appendix for an example of how banks communicate their actual mortgage interest rates. To exploit the unique features caused by the Single Interest Rate policy, I only use mortgages that originated in 2013 or later.

In my analysis I do not take into account mortgages that are insured under the National Mortgage Guarantee (*Nationale Hypotheekgarantie*, NHG). This group of mortgages is a very safe investment for banks since a government-backed fund guarantees to compensate banks in the case that an NHG mortgage defaults. Therefore, NHG mortgages comprise a distinct risk bracket with a discounted interest rate. Since this rate does not depend on the LTV level, households in this category do not face a notched budget set, rendering it incompatible with bunching analysis.

4.2 Data

The data used in this analysis come from the Loan Level Data (LLD) database of De Nederlandsche Bank (DNB), which is the Dutch central bank. DNB collects mortgage level data from mortgage suppliers quarterly. Despite the fact that banks are not obliged by law to hand over these details, the LLD cover roughly 70-80% of mortgage debt outstanding in the Netherlands. The dataset includes many details on the mortgage contract, such as origination details (date of origination, provider, etc.), loan amount, LTV at origination, the type of mortgage (annuities, linear etc.), interest rate type (fixed, flexible etc.), and the duration of the mortgage (time between origination and maturity). Moreover, the LLD provide some borrower characteristics, such as income, age, whether a household has one or two breadwinners, and whether a household falls under the NHG. Another important – and quite fortunate – attribute of the LLD is that it reports LTV levels net of origination fees.[18] This is important because the mortgage interest rate depends on the LTV net of origination fees. For most observations, the banks reported the 'official' LTV at the mortgage origination to DNB. In some instances, banks failed to disclose this static. For these observations, the LTV was calculated manually by DNB.

It should be noted that banks usually set different interest rates for the different mort-

[16] This page contains the text of the legislation (`https://tinyurl.com/yaeu2e6h`).

[17] The text of the article reads (in Dutch): "*Een aanbieder van een hypothecair krediet offreert voor een consument die voornemens is een overeenkomst inzake een hypothecair krediet aan te gaan dezelfde debetrentevoet bij dezelfde rentevastperiode als voor een consument aan wie op dat moment een aanbod wordt gedaan voor de komende rentevastperiode bij een vergelijkbaar risicoprofiel.*"

[18] For the data used by Best et al. (2020) the origination fees were rolled into the LTV calculation, but without affecting the interest rate calculation. As a result, these authors had to correct all reported LTVs for origination fees and drop data for which origination fee information could not be retrieved.

gage and interest types. For example, the interest rate of a mortgage with a given LTV might be higher when the repayment scheme is linear instead of in annuities. However, the percentage point increment in the interest rate at the relevant LTV thresholds is the same across all types for the banks that I study. Thus, the LTV does not dissimilarly affect interest rates of different mortgage and interest types.

Furthermore, I focus on newly created mortgages, thus disregarding households that refinance or reappraise their dwelling, or partly pay off their mortgage. Doing so avoids that the sample contains households whose LTV drops below the threshold over time.

From the universe of mortgages in the LLD, I have a random sample of 100,000 mortgages at my disposal. Deleting the loans that originated before 2013 – to be able to exploit the Single Interest Rate policy – results in a data set that has shrunk drastically. An additional benefit of only using observations after 2012 is that this eliminates measurement error. For some mortgages that were sold before 2012 there is a discrepancy between the origination and valuation date of mortgages. For these observations, DNB used an imputed number to proxy the value a dwelling had at the mortgage origination date. On average, this imputed value is be correct. However, since bunching analysis requires the precise LTV that a household faced at origination, deleting the observations from this period is beneficial.

As a result of my using only post-2013 mortgages, I do not have enough observations in my sample to be able to plot the LTV density for all banks. Therefore, I limit the analysis to three of the biggest banks: ABN AMRO, Rabobank and SNS. Of these banks, I have roughly 15,000 post-2013 observations in total. In accordance with the fact that many Dutch dwellings are fully leveraged, most of the 15,000 post-2013 mortgages of the three banks (63%) are high up in the LTV distribution. This means that many mortgages in the sample are not 'bound' by the interest notches, leaving some 10,000 observations for the bunching analysis.

4.3 Counterfactual distribution and behavioural response

The theoretical description of the bunching estimator as developed earlier in this paper explicitly relied on a counterfactual distribution. This would be the distribution had there not been notches in the mortgage interest rate schedule. The standard approach to obtain a counterfactual in the bunching literature is to fit a 'flexible polynomial' through the observed distribution, leaving out an interval around the notch. I will closely follow the seminal paper by Kleven and Waseem (2013) who extend and further formalise the counterfactual estimation from Chetty et al. (2011). My specification is comparable to the one by Best and Kleven (2018), whose paper is representative of the bunching literature in the sense that it uses the most common and modern methodology to estimate counterfactuals in a 'notched' setting. My counterfactual estimation is, however, not related to the one by Best et al. (2020), despite the fact that their paper studies a similar institutional setting.[19]

[19] Although the institutional setting of my paper is similar to the one in Best et al. (2020), who use notches in the interest rate schedule in the UK retrieve a structural estimate of the elasticity of intertemporal

My methodology is as follows. I group all mortgages in LTV bins that span 0.25 percentage points and fit a flexible polynomial through the obtained distribution, whilst excluding observations in a range around the LTV notches. First, I specify the regression

$$n_j = \sum_{i=0}^{q} \beta_i (\ell_j)^i + \sum_{i=\ell_L}^{\ell_H} \gamma_i \cdot \mathbf{I}\left[\ell_j = i\right] + \eta_j, \tag{16}$$

where n_j is the number of mortgages in LTV bin j, ℓ_j is the LTV ratio in bin j, β_i is the 'coefficient', and q is the order of the polynomial. The second term in equation (16) assigns dummies (γ_i) to the bin-counts in the region $[\ell_L, \ell_H]$ that is distorted from bunching behaviour as a result of the notch; $\mathbf{I}[\cdot]$ functions as an indicator function. Lastly, η_i is the residual term that captures variation due to misspecification of the density.

To obtain the counterfactual amount of mortgages per LTV bin I take the predicted values $\hat{n}_j$ without the effect of the bins in the excluded range, such that the counterfactual density can be defined as

$$\hat{f}_0(\ell) \equiv \hat{n}_j = \sum_{i=0}^{q} \hat{\beta}_i (\ell_j)^i. \tag{17}$$

Now, comparing the actual and counterfactual distributions reveals the excess bunching and missing mass. Bunching, $\hat{B}$, is defined as the amount by which the observed distribution exceeds the counterfactual in a region before notch ℓ^τ, such that

$$\hat{B} = \sum_{i=\ell_L}^{\ell^\tau} \left(n_j - \hat{n}_j\right). \tag{18}$$

This specification allows for 'overshooting' by households trying to bunch under the threshold. Missing mass $\hat{M}$ is given by the difference between the observed and counterfactual bin counts in the other part of the excluded region,

$$\hat{M} = \sum_{i>\ell^\tau}^{\ell_H} \left(\hat{n}_j - n_j\right). \tag{19}$$

The excluded region $[\ell_L, \ell_H]$ spans the entire range that is affected by the behavioural response. Since the bunching response is sharp, the lower bound ℓ_L can be determined unambiguously upon visual inspection. However, it is harder to establish the upper bound ℓ_H. I follow the approach by Kleven and Waseem (2013) that – conditional on the absence of extensive margin effects – requires the condition that the excess bunching mass is equal to the missing mass above the notch. This means that I estimate iterations of (16) together

substitution (EIS), the methodologies to estimate the counterfactual are different. These authors explicitly look at refinancers, which enables them to exploit a panel structure (i.e., they observe an household before and after the refinancing decision). The observed LTV just before refinancing is labelled as the 'passive LTV'. They correct this passive LTV with the average extra mortgage amount that is taken out when households refinance and use this as their counterfactual. This approach is facilitated by a unique institutional setting where households refinance every five years, but is not applicable to the Dutch mortgage market. Note that the data used by Best et al. (2015, 2020) is used for a different type of analysis in the work by Cloyne, Huber, Ilzetzki, and Kleven (2017, 2019).

with ℓ_H, ensuring that the missing mass $\hat{M}$ equates the bunching mass $\hat{B}$. I start with a low value of the upper bound, $\ell_H^0 \approx \ell^\tau$, to retrieve the corresponding counterfactual bin counts $\hat{n}_j^0$, where it holds that $\hat{M}^0 \ll \hat{B}^0$. Then, in iterations, I expand the upper bound with small steps until the corresponding counterfactual is such that $\hat{M}^k = \hat{B}^k$. This yields the upper bound $\ell_H = \hat{\ell}_H^k$.[20] Furthermore, I check (for a range of polynomial degrees q) that ℓ_H is smaller than notches higher up in the LTV distribution. This is to ensure that there is no bias due to bunchers jumping more than one notch – something that the empirical bunching methodology is not robust against.

The processes described above provides all the necessary ingredients to back out the (average) behavioural response $\Delta\ell$. As described in section 3.4, there are two ways to do this. The first method, which I will refer to as the 'Bunching-hole approach', requires the assumption that the density is constant throughout the entire range of bunching individuals. The second method, labelled 'Convergence approach', drops this assumption.[21]

Throughout this section, standard errors are obtained using the bootstrap procedure pioneered by Kleven and Waseem (2013). The idea is to generate a large amount of LTV distributions by randomly resampling the residuals from equation (16). Then, the standard error for each variable of interest is given by the standard deviation of that variable in the randomly generated distribution. In this paper, I randomly rerun each estimate 100 times – as is common in the literature.

4.3.1 Bunching-hole approach

As shown in the equations (14) and (15) in the theoretical specification, dividing the bunching mass by the density at the threshold is in approximation equal to the (mean) behavioural response. I label this approximation $\Delta\ell_b$, such that $\Delta\ell_b \equiv \hat{B}/\hat{f}_0(\ell^\tau)$. Recall that $\hat{B}$ allows for overshooting by summing bunching on a small interval below the notch. Because of this, I choose to define the counterfactual density at the notch as the average counterfactual bin count in the 'overshooting region'. This means that I divide the sum of counterfactual bin counts in this region by the number of bins, so that $\hat{f}_0(\ell^\tau) = \frac{1}{(\ell^\tau - \ell_L)/0.25+1} \sum_{i=\ell_L}^{\ell^\tau} \hat{n}_j$.

My estimates of $\Delta\ell_b$ are robust against different specifications of the counterfactual density, but I maintain the 'average counterfactual' because this method seems to be good practice in the recent literature (see for instance Best & Kleven, 2018). The bins that I use in the analysis span 0.25 percentage points. To sum this all up,

$$\Delta\ell_b \equiv \frac{\hat{B}}{\sum_{i=\ell_L}^{\ell^\tau} \hat{n}_j} \left(\frac{\ell^\tau - \ell_L}{0.25} + 1 \right). \tag{20}$$

There are some explicit advantages and disadvantages to this bunching-hole approach. The main advantage of this method is that $\Delta\ell_b$ does not rely on an an estimate of the

[20] Diamond and Persson (2016) have developed a framework that allows extends the estimation of the upper bound to the lower bound – given a parametric assumption on the functional form of the counterfactual. I do not follow their approach, since visual inspection suffices to find the lower bound in my case.

[21] The 'Bunching-hole' and 'convergence' approaches are Kleven and Waseem (2013) terminology.

counterfactual *above* the notch. This is a good thing for two reasons. First, as explained by Kleven and Waseem (2013), the estimation of the counterfactual results in a 'partial counterfactual' that gets rid of the intensive margin responses induced by the notch. However, the counterfactual cannot correct for responses on the extensive margin (e.g., some households might switch to a bank that has set slightly different interest thresholds). Since the notch only affects individuals above the LTV threshold, the counterfactual below the threshold is unlikely to be biased by extensive margin effects. Of course, it could be the case that bunching $\hat{B}$ is biased, due to 'newcomers' who switch to a sub-notch LTV with the bank under analysis. This would not be picked up by the counterfactual bin count $\sum_{i=\ell_L}^{\ell_{j^\tau}} \hat{n}_j$, potentially leading to an overestimation of $\hat{B}$ – and correspondingly an overestimation of $\Delta\ell_b$. With the data at my disposal it is not possible to check how strong these 'between-bank' extensive margin responses are. In theory, it is possible that some households in the bunching region came from different banks, whereas households in the missing mass region moved to other banks. However, as long as both biases are of equal size (i.e., the group that moves to other banks is as big as the group that comes from another bank) both biases would be offsetting.

The second advantage of the bunching-hole approach has to do with the fact that it is hard to determine the correct location of the upper bound of the excluded region $[\ell_L, \ell_H]$, since the hole in the distribution is diffuse (and large). Picking a wrong upper bound might lead to a positively or negatively biased estimate of the counterfactual above the notch. Furthermore, since missing mass $\hat{M}$ is sensitive to the polynomial order of the regression (Best & Kleven, 2018), polynomial misspecification might lead to a similar bias. Thus, by disregarding the counterfactual above the notch, equation (20) is robust against biases in the counterfactual that are induced by extensive margin responses and misspecification of the counterfactual.

The main disadvantage of this approach is that it assumes that the counterfactual is a constant function. Although $f_0(\ell)$ is smooth and relatively stable, it is by no means locally constant.

4.3.2 Convergence approach

The convergence method is similar to the upper bound estimation described towards the end of section 4.3. Given the specification of the counterfactual, the convergence approach retrieves the behavioural response $\Delta\ell_c$ by increasing it with small increments until $\hat{M} = \hat{B}$. This means that I start with a low value of the behavioural response $\Delta\ell_c^0 \approx \ell^\tau$, so that $\hat{M}^0 \ll \hat{B}^0$. Then, in iterations, I increase the behavioural response until I find $\Delta\ell_c^k$, for which it holds that $\hat{M}^k = \hat{B}^k$. The advantage of this approach is that it does not require the same functional form assumption as the bunching-hole approach (constant counterfactual). However, there can be bias due to a misspecification of the counterfactual above the notch, most likely as a result of responses on the extensive margin. Be that as it may, ℓ_c can be regarded as an lower bound for the behavioural response of the marginal buncher (Kleven & Waseem, 2013).

4.3.3 Pooled notch

First, I analyse all notches separately: for each specific notch I estimate the behavioural response and calculate the reduced form elasticities. Then, I pool the loan data into a single average notch to be able to consider all notches together. To do so, I normalise LTVs by defining the LTV of each mortgage i relative to notch τ as $\text{LTV}_{i\tau} = \text{LTV}_i - \tau$. This means that a mortgage with an LTV of 87% has a normalised LTV of $\text{LTV}_{i85} = 2$ with respect to the 85% notch and a normalised LTV of $\text{LTV}_{i90} = -3$ with respect to the 90% notch. As a result, I can simply set the pooled notch at zero and stack the observed and counterfactual distributions of the mortgages around this, allowing me to calculate the behavioural response across the six notches. Since I stack the already estimated notch specific counterfactuals, I do not need to determine an excluded region. Lastly, due to the fact that interest jumps differ across banks and notches, I use weighted averages of the interest rate, interest rate jump and threshold levels in the reduced-form elasticity calculations for the pooled notch.

5 Estimates

Having established the theoretical and empirical framework, presenting the estimates is straightforward. Figure 4 shows the observed and counterfactual loan-to-value ratio (LTV) distributions when the notches are pooled, and provides the estimated behavioural responses. Figure 5 shows these same estimates for each individual notch, using my favourite specification of the counterfactual. For a further reference, see the Figures B.3 – B.8 in the appendix. These graphs provide a robustness check of the effects found by estimating the counterfactual under a different polynomial order.

Figure 4: Bunching estimates when the notches are pooled

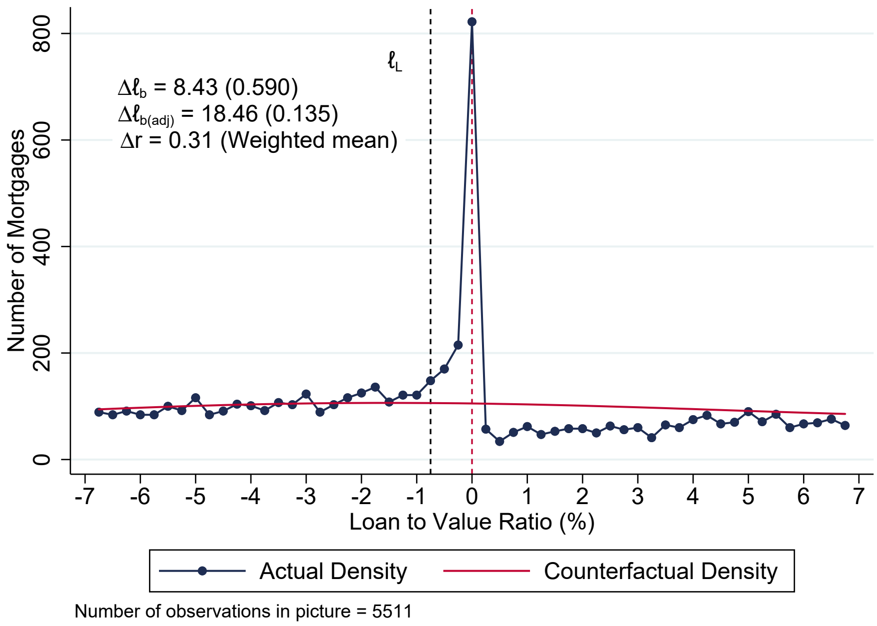

Notes: This graph displays the observed and counterfactual distributions of the LTV when pooling the notches. The behavioural responses are denoted in percentage points, with standard errors obtained from the bootstrap procedure in parentheses. The dotted black line segment marks the lower bound on the excluded region. The applied binwidth is 0.25 percentage points.

Figure 4 and 5 already contain some of the essential empirical moments needed for the reduced-form elasticity calculations. $\Delta\ell_c$ is the behavioural response obtained using the convergence method, whereas $\Delta\ell_b$ and $\Delta\ell_{b(adj.)}$ are obtained using the bunching-hole method where the latter is adjusted for optimisation frictions, as described in Section 3.5. Moreover, the figures contain information on interest jump Δr, polynomial order q, upper bound ℓ_H (as described in Section 4.3), as well as the total number of mortgages depicted in a specific graph.

When looking at the figures the following findings emerge. First of all, every graph displays sharp bunching that is on average 3.5 times bigger than the estimated counterfactual distribution at the notch. Secondly, in all settings there is a clearly visible hole between the observed and counterfactual distributions above the notch. The distance between notches

Figure 5: Bunching estimates by bank and notch

Panel A: *Rabobank – 67.5 notch*

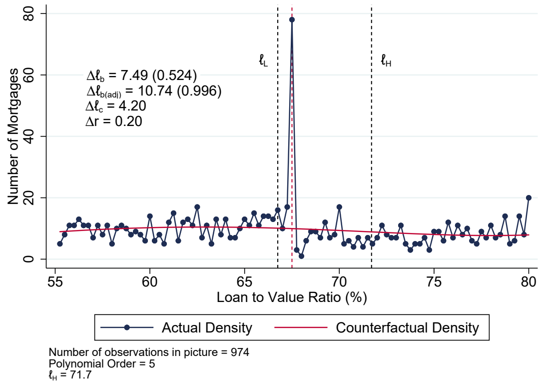

Panel B: *Rabobank – 90 notch*

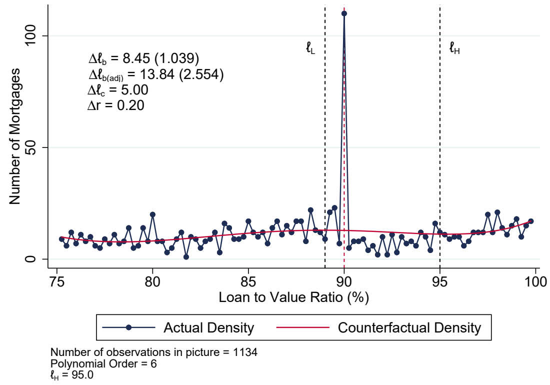

Panel C: *ABN AMRO – 65 notch*

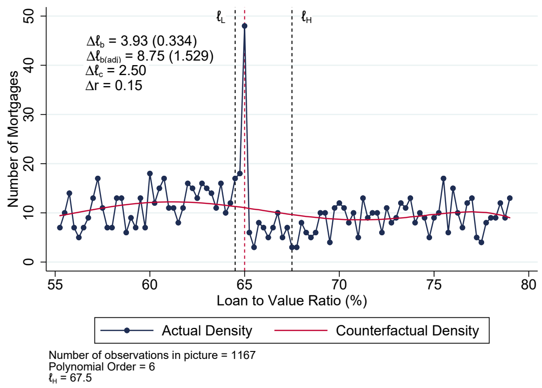

Panel D: *ABN AMRO – 85 notch*

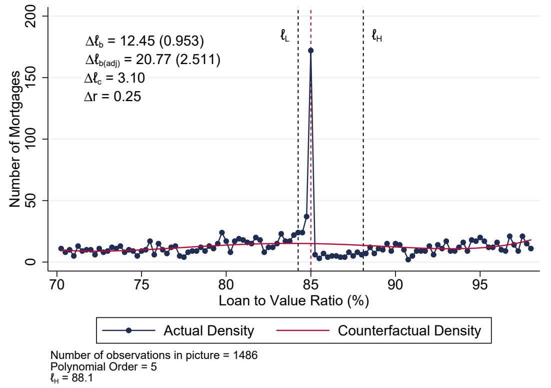

Panel E: *SNS – 66 notch*

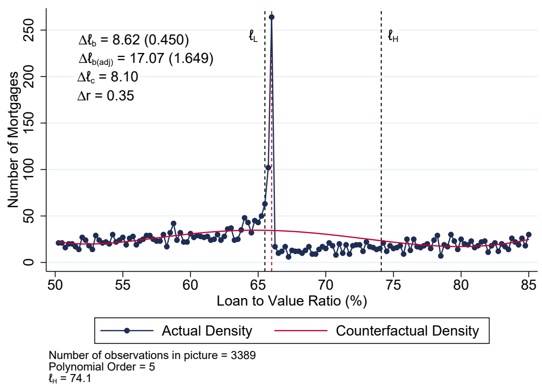

Panel F: *SNS – 88 notch*

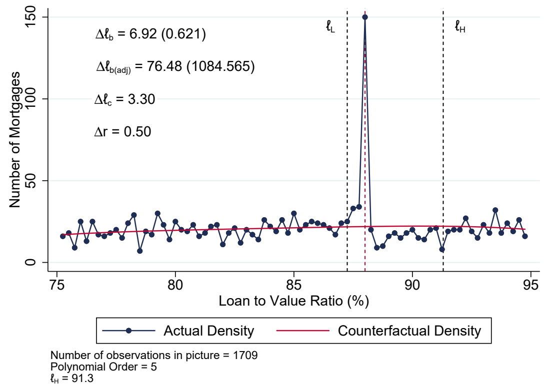

Notes: These graphs display the observed and counterfactual distributions of the LTV for the three banks around their respective LTV thresholds. The behavioural responses are denoted in percentage points, with standard errors obtained from the bootstrap procedure in parentheses. The two dotted black line segments mark the bounds on the excluded region $[\ell_L, \ell_H]$. The applied binwidth is 0.25 percentage points; the polynomial order used to estimate the counterfactual is provided below the figure.

and the end of this hole varies by bank and LTV threshold. In all cases this missing mass ends well before other notches, which implies that households do not move down more than one notch when they bunch. Thirdly, behavioural responses are generally larger for notches at higher LTV levels. Furthermore, for no notch setting the observed distribution is zero in a 'hole region'. As discussed in Section 3.5, this likely reflects the existence of optimisation frictions such as misperception, inattention or credit and liquidity constraints. However, the correction used for this is an imprecise measure. As such, the estimated behavioural response that results from this should be regarded as an extreme upper bound. It is not realistic that more than fifty percent of households cannot bunch due to optimisation frictions. A possible explanation of this high fraction is that some dwellings are financed with a mortgage that has a flexible interest rate. When these mortgages are partly paid off in the future, their LTVs and corresponding interest rates will drop below the threshold. Households with a mortgage of this type thus have a smaller incentive to bunch, explaining the high fraction of 'passive' households.

Lastly, a note should be made on the estimate retrieved using the convergence approach ($\Delta\ell_c$). As explained in section 4.3.2, this estimate is prone to biases due to misspecification of the counterfactual above the notch. Moreover, visual inspection of the graphs in Figure 5 shows that the observed distribution has a high variance for the banks with a relatively small number of observations (Rabobank & ABN AMRO). Most likely, this volatility of the observed distribution is due the limited sample size. As a result, there could be a bias in the missing mass calculation. Because of this, I will take the 'bunching-hole' estimate as the best measure of the behavioural response of the average marginal buncher. This choice is in line with the recent literature, see for instance the work by DeFusco and Paciorek (2017) and Best and Kleven (2018).

Table 2 presents an overview of the behavioural responses, mean interest rates at the notch as well as the size of the interest jump. Also, this table presents the reduced-form estimates obtained from the formulae in Section 3.3. My preferred bunching statics are the non-adjusted bunching-hole estimates, which are always on the middle row of a panel. The estimated semi-elasticities range from -0.011 to -0.075, with a pooled estimate of -0.035 (i.e., 1.1%, 7.5% and 3.5%). These (absolute values of the) semi-elasticities should be interpreted as the decrease in the mortgage balance in percentages given a one percentage point increase in the interest rate.

How should the economic magnitude of these estimates be understood? In absolute terms, the effects found are sizeable. For example, the pooled semi-elasticity relates a 1 percentage point increase of the mortgage interest rate to a mortgage reduction of 3.5%. Naturally, caution is advised when using these estimates outside of the context they were estimated in. Since the results are reduced-form estimates rather than the outcome of a structural model they might be unique to the exact setting they were estimated in. For example, different LTV levels, other baseline interest rates or different government polices (e.g., a higher/lower mortgage interest deduction rate) might have led to different figures. However, it is still interesting to see that the estimates from Table 2 are slightly

Table 2: Reduced-form estimates of mortgage demand

	Notch						
	Rabobank		ABN AMRO		SNS		All
Statistic	**67.5**	**90**	**65**	**85**	**66**	**88**	**Pooled**
Panel A: Behavioural Responses							
$\Delta\ell_c$	4.20	5.00	2.50	3.10	8.10	3.30	-
$\Delta\ell_b$	7.49	8.45	3.93	12.45	8.62	6.92	8.43
	(0.524)	(1.039)	(0.334)	(0.953)	(0.450)	(0.621)	(0.590)
$\Delta\ell_{b(\text{adj})}$	10.74	13.84	8.75	20.77	17.07	76.48	18.46
	(0.996)	(2.554)	(1.529)	(2.511)	(1.649)	(1085)	(0.135)
Panel B: Interest Rates							
r (%)	2.36	2.67	2.22	2.77	3.88	4.24	2.97
	(0.070)	(0.051)	(0.079)	(0.069)	(0.079)	(0.120)	(0.056)
Δr (pp)	0.20	0.20	0.15	0.25	0.35	0.50	0.31
Panel C: Reduced-form semi-elasticities (ε_s)							
Conv.	0.018	0.015	0.009	0.005	0.038	0.003	-
Bunch.-h	0.055	0.040	0.023	0.075	0.043	0.011	0.035
	(0.007)	(0.010)	(0.004)	(0.011)	(0.004)	(0.002)	(0.005)
B.-h (adj.)	0.109	0.102	0.106	0.192	0.152	0.808	0.149
	(0.019)	(0.037)	(0.039)	(0.043)	(0.026)	(25.35)	(0.033)
Panel D: Reduced-form elasticities (ε)							
Conv.	0.105	0.095	0.060	0.051	0.272	0.049	-
Bunch.-h	0.242	0.201	0.111	0.354	0.298	0.127	0.213
	(0.025)	(0.037)	(0.014)	(0.041)	(0.023)	(0.016)	(0.022)
B.-h (adj.)	0.417	0.427	0.371	0.776	0.849	4.294	0.683
	(0.060)	(0.128)	(0.110)	(0.149)	(0.127)	(119.5)	(0.098)

Notes: This table contains the reduced-form semi-elasticities (ε_s) and elasticities (ε) for each notch of the three banks, as well as for a pool of all notches. Panel A displays the behavioural responses for the conventional, bunching-hole and bunching-hole with optimisation frictions approaches. Panel B contains the mean interest rate at the notch (standard errors in parentheses) alongside the corresponding interest rate jump. The other panels show the reduced-form elasticities. The displayed standard errors were calculated using the aforementioned bootstrap procedure, where it should be noted that it is not possible to retrieve standard errors for the convergence statics. For the elasticity calculation of the pooled notches I used the average threshold value (76.8) as ℓ^τ.

higher than those found in the existing literature. Fuster and Zafar (2015) execute a 'strategic survey' in which they evaluate households' willingness to pay for a dwelling under hypothetical financing schemes. These authors find small behavioural responses and report semi-elasticity estimates between 1.8% and 2.3%. In a paper methodologically related to mine, DeFusco and Paciorek (2017) exploit the maximum loan size upheld by two banks in the United States. The authors report semi-elasticities ranging from 1.5% to 2%.

Lastly, it should be noted that with the limited sample at my disposal it was not possible to credibly identify heterogeneous treatment effects. Bunching analysis can only be properly executed on large data sets. In comparison with other bunching papers, the number of observations that I use is small. Breaking this sample down in smaller subsets to retrieve heterogeneous effects simply does not work – which was also the reason for my analysing only the three largest banks. On the road to further research, analysing the full Loan Level Data (LLD) will be highly beneficial. Not only will this allow for the investigation of the effects for the other mortgage providers, but will also enable to distinguish between different mortgage types (e.g., fixed/flexible interest rates, linear/annuity repayment). Moreover, looking at the full LLD would allow for the analysis of heterogeneous treatment effects at the level of household characteristics such as family size, age, loan-to-income ratio and others. Moreover, it would make it possible to distinguish between dwelling characteristics such as its size, price and location.

6 Conclusion

This paper set out to estimate the relation between interest rates and mortgage debt. To do so, this paper exploits a quasi-experiment in the Dutch mortgage market. Interest rates schedules feature discrete jumps at several loan-to-value ratio (LTV) levels. These interest rate notches induce households to bunch at LTV levels just below a notch. The paper uses this behaviour to estimate reduced-form elasticities of mortgage demand. This relation is beneficial to policy makers who want to know the effect that interest rate policies have on the size and composition of mortgage debt. Current examples of this are the ECB's monetary policy and reductions in the mortgage interest deduction (MID) rate.

This estimation is not perfect. First of all, it depends on the assumption that households can only influence the numerator of the fraction $\frac{\text{mortgage size}}{\text{house value}}$ to end up below a notch – the market value of a dwelling is 'given'. However, there is the theoretical possibility that appraisers overestimate the market value of a house to lower the LTV. In that case, a part of the observed bunching would not be the result of reductions in loan size. This would lead to an overestimation of the behavioural response and the elasticity of mortgage demand. Nevertheless, this problem remains a theoretical possibility. There is no proof for it – but also no evidence that disproves the theory.

Moreover, the bunching methodology cannot correct for extensive margin responses. Since not all banks have notches at the same LTV levels, it is possible that households move to banks that have an interest jump at a more favourable LTV threshold. If a banks attracts more of these households than it loses to other banks, the bunching mass relative to the missing mass is overstated, which leads to an upward bias in the estimates. Vice versa, if a bank 'loses' more households, the missing mass is overstated. If both effects are equal in size, the biases are offsetting. It is not unreasonable to assume that this is the case. However, with the data used in this paper it is not possible to check to what extent these extensive margin responses play a role.

Furthermore, it is likely that some households do not optimise due to their being constrained by optimisation frictions, such as credit and liquidity constraints. Since the correction for this is an imprecise measure, it is not possible to correct for these frictions.

Future research should address these three points. Last but not least, the following has to be stressed. The analysis of this paper could be greatly improved by rerunning the study using the universe of mortgage loans in the Netherlands. A bigger sample will have a less volatile observed distribution, thus improving the credibility of the counterfactual estimates. In addition, using all mortgages mitigates the potential extensive margin bias described above. If all banks and all customers are included, this bias does not apply to the pooled estimate. Most importantly, having a bigger data set would allow for the analysis of heterogeneous treatment effects.

References

Almunia, M., & Lopez-Rodriguez, D. (2018). Under the radar: The effects of monitoring firms on tax compliance. *American Economic Journal: Economic Policy*, *10*(1), 1-38. Link to paper.

Attanasio, O. P., Goldberg, P. K., & Kyriazidou, E. (2008). Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review*, *49*(2), 401-436. Link to paper.

Best, M., Cloyne, J., Ilzetzki, E., & Kleven, H. J. (2015). *Interest rates, debt and intertemporal allocation: Evidence from notched mortgage contracts in the United Kingdom* (Staff Working Paper No. 543). Bank of England. Link to paper.

Best, M., Cloyne, J., Ilzetzki, E., & Kleven, H. J. (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies*, *87*(2), 656-690. Link to paper.

Best, M., & Kleven, H. J. (2018). Housing market responses to transaction taxes: Evidence from notches and stimulus in the United Kingdom. *The Review of Economic Studies*, *85*(1), 157-193. Link to paper.

Bosch, N., Jongen, E., Leenders, W., & Möhlmann, J. (2019). Non-bunching at kinks and notches in cash transfers in the Netherlands. *International Tax and Public Finance*. Link to paper.

Brehm, M., Imberman, S. A., & Lovenheim, M. F. (2017). Achievement effects of individual performance incentives in a teacher merit pay tournament. *Labour Economics*, *44*, 133 - 150. Link to paper.

Brown, K. M. (2013). The link between pensions and retirement timing: Lessons from California teachers. *Journal of Public Economics*, *98*, 1 - 14. Link to paper.

Brueckner, J. K. (1994). The demand for mortgage debt: Some basic results. *Journal of Housing Economics*, *3*(4), 251 - 262. Link to paper.

Card, D., Lee, D. S., Pei, Z., & Weber, A. (2015). Inference on causal effects in a generalized regression kink design. *Econometrica*, *83*(6), 2453-2483. Link to paper.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*. Link to paper.

Chetty, R., Friedman, J. N., Olsen, T., & Pistaferri, L. (2011). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records. *The Quarterly Journal of Economics*, *126*(2), 749-804. Link to paper.

Cloyne, J., Huber, K., Ilzetzki, E., & Kleven, H. (2017). *The effect of house prices on household borrowing: A new approach* (Working Paper No. 23861). National Bureau of Economic Research. Link to paper.

Cloyne, J., Huber, K., Ilzetzki, E., & Kleven, H. (2019). The effect of house prices on household borrowing: A new approach. *American Economic Review*, *109*(6), 2104-36. Link to paper.

DeFusco, A. A., & Paciorek, A. (2017). The interest rate elasticity of mortgage demand:

Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy*, *9*(1), 210-40. Link to paper.

De Mooij, R., Van Ewijk, C., & Jacobs, B. (2006). Doelmatigheidswinst van minder hypotheekrenteaftrek. *Economisch-Statistische Berichten*, 292–295.

Diamond, R., & Persson, P. (2016). *The long-term consequences of teacher discretion in grading of high-stakes tests* (Working Paper No. 22207). National Bureau of Economic Research. Link to paper.

Einav, L., Finkelstein, A., & Schrimpf, P. (2015). The response of drug expenditure to nonlinear contract design: Evidence from Medicare Part D. *The Quarterly Journal of Economics*, *130*(2), 841-899. Link to paper.

Einav, L., Finkelstein, A., & Schrimpf, P. (2019). Bunching at the kink: Implications for spending responses to health insurance contracts. *Journal of Public Economics*, *171*, 117 - 130. Link to paper.

Finkelstein, A., Hendren, N., & Shepard, M. (2019). Subsidizing health insurance for low-income adults: Evidence from Massachusetts. *American Economic Review*, *109*(4), 1530-67. Link to paper.

Follain, J. R., & Dunsky, R. M. (1997). The demand for mortgage debt and the income tax. *Journal of Housing research*, *8*(2), 155-188. Link to paper.

Fuster, A., & Zafar, B. (2015). The sensitivity of housing demand to financing conditions: evidence from a survey. *FRB of New York Staff Report*(702). Link to paper.

Goupille-Lebret, J., & Infante, J. (2018). Behavioral responses to inheritance tax: Evidence from notches in France. *Journal of Public Economics*, *168*, 21 - 34. Link to paper.

Hoekstra, W. (2019). *Belastingplan 2020* (Kamerstukken 35 026 No. 2).

Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, *142*(2), 615 - 635. Link to paper.

Ito, K., & Sallee, J. M. (2018). The economics of attribute-based regulation: Theory and evidence from fuel economy standards. *The Review of Economics and Statistics*, *100*(2), 319-336. Link to paper.

Jakobsen, K., Jakobsen, K., Kleven, H., & Zucman, G. (2018). *Wealth taxation and wealth accumulation: Theory and evidence from Denmark* (Working Paper No. 24371). National Bureau of Economic Research. Link to paper.

Jappelli, T., & Pistaferri, L. (2007). Do people respond to tax incentives? An analysis of the Italian reform of the deductibility of home mortgage interests. *European Economic Review*, *51*(2), 247 - 271. Link to paper.

Jordà, O., Schularick, M., & Taylor, A. M. (2015). Betting the house. *Journal of International Economics*, *96*(1), 2-18. Link to paper.

Karlan, D. S., & Zinman, J. (2008). Credit elasticities in less-developed economies: Implications for microfinance. *American Economic Review*, *98*(3), 1040-68. Link to paper.

Keys, B. J., & Wang, J. (2019). Minimum payments and debt paydown in consumer credit cards. *Journal of Financial Economics*, *131*(3), 528 - 548. Link to paper.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, *8*(1), 435-464. Link to paper.

Kleven, H. J., & Waseem, M. (2013). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. *The Quarterly Journal of Economics*, *128*(2), 669-723. Link to paper.

Kopczuk, W., & Munroe, D. (2015). Mansion tax: The effect of transfer taxes on the residential real estate market. *American Economic Journal: Economic Policy*, *7*(2), 214-57. Link to paper.

Manoli, D., & Weber, A. (2016). Nonparametric evidence on the effects of financial incentives on retirement decisions. *American Economic Journal: Economic Policy*, *8*(4), 160-82. Link to paper.

Martins, N. C., & Villanueva, E. (2006). The impact of mortgage interest-rate subsidies on household borrowing. *Journal of Public Economics*, *90*(8), 1601 - 1623. Link to paper.

Poterba, J. M. (1984). Tax subsidies to owner-occupied housing: An asset-market approach. *The Quarterly Journal of Economics*, *99*(4), 729 - 752. Link to paper.

Poterba, J. M., & Sinai, T. (2008). Tax expenditures for owner-occupied housing: Deductions for property taxes and mortgage interest and the exclusion of imputed rental income. *American Economic Review*, *98*(2), 84-89. Link to paper.

Ruh, P., & Staubli, S. (2019). Financial incentives and earnings of disability insurance recipients: Evidence from a notch design. *American Economic Journal: Economic Policy*, *11*(2), 269-300. Link to paper.

Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, *2*(3), 180-212. Link to paper.

Sallee, J. M., & Slemrod, J. (2012). Car notches: Strategic automaker responses to fuel economy policy. *Journal of Public Economics*, *96*(11), 981 - 999. Link to paper.

Søgaard, J. E. (2019). Labor supply and optimization frictions: Evidence from the Danish student labor market. *Journal of Public Economics*, *173*, 125 - 138. Link to paper.

Van Ewijk, C., Jacobs, B., & De Mooij, R. (2007). Welfare effects of fiscal subsidies on home ownership in the Netherlands. *De Economist*, *155*(3), 323–336. Link to paper.

Van Ewijk, C., & Lejour, A. (2019). Ontwerp voor een beter belastingstelsel. In S. Cnossen & B. Jacobs (Eds.), (p. 32 - 45). Ten Brink, Meppel.

Appendices

Appendix A: A model of LTV choice

A.1 Optimal LTV choice

The household's optimisation problem is given by:

$$\max_{c_1,c_2}\left\{U(c_1,c_2)\equiv u(c_1)+\delta u(c_2)=\frac{1}{1-\sigma}c_1^{1-\sigma}+\delta\frac{1}{1-\sigma}c_2^{1-\sigma}\right\} \tag{A.1}$$

$$\text{s.t.}\quad c_1=W_0-(1-\ell)p \tag{A.2}$$

$$c_2=p[1-\ell(1+r)], \tag{A.3}$$

where $\delta\in(0,1)$ discounts future utility, W_0 is the sum of period-one income and assets, r is the mortgage interest rate, and ℓ is the loan-to-value ratio (LTV). Solving this problem is straightforward. Substitute the equations (A.2) and (A.3) in equation (A.1). This incorporates the household budget constraints in the objective function and expresses this function in terms of ℓ, which households now pick in order to optimise their total utility:

$$\max_{\ell}\left\{V\equiv\frac{1}{1-\sigma}[W_0-(1-\ell)p]^{1-\sigma}+\delta\frac{1}{1-\sigma}[p(1-\ell(1+r))]^{1-\sigma}\right\}. \tag{A.4}$$

Taking the partial derivative of V in (A.4) with respect to the LTV and setting this equal to zero yields an expression that can be solved for the optimal ℓ in equation (5):

$$\frac{\partial V}{\partial \ell}=[W_0-(1-\ell)p]^{-\sigma}p-\delta p(1+r)[p(1-\ell(1+r))]^{-\sigma}=0$$

$$\Leftrightarrow W_0-(1-\ell)p=\frac{p(1-\ell(1+r))}{(\delta(1+r))^{\frac{1}{\sigma}}}$$

$$\Leftrightarrow \ell^*=\frac{1+(1-\frac{W_0}{p})\,(\delta(1+r))^{\frac{1}{\sigma}}}{(1+r)+(\delta(1+r))^{\frac{1}{\sigma}}}.$$

A.2 Implicit definition of the elasticity

The marginal buncher is indifferent between the LTV levels ℓ^τ and ℓ^I, where $\ell^I\in(\ell^\tau,\ell+\Delta\ell)$. These two points correspond to the utility levels $\mathcal{U}^N$ and $\mathcal{U}^I$, respectively. The utility level at the LTV threshold ℓ^τ is given by

$$\mathcal{U}^N=\frac{1}{1-\sigma}\Big(\underbrace{W_0-(1-\ell^\tau)p}_{c_1}\Big)^{1-\sigma}+\delta\frac{1}{1-\sigma}\Big(\underbrace{p[1-\ell^\tau(1+r)]}_{c_2}\Big)^{1-\sigma}. \tag{A.5}$$

The utility level at the 'indifference point' is

$$\mathcal{U}^I = \frac{1}{1-\sigma}\Bigg(\underbrace{W_0 - (1-\ell^I)p}_{c_1}\Bigg)^{1-\sigma} + \delta\frac{1}{1-\sigma}\Bigg(\underbrace{p[1-\ell^I(1+r+\Delta r)]}_{c_2^n}\Bigg)^{1-\sigma}, \tag{A.6}$$

where it should be noted that

$$\ell^I = \frac{1+\left(1-\frac{W_0}{p}\right)(\delta(1+r+\Delta r))^{\frac{1}{\sigma}}}{(1+r+\Delta r)+(\delta(1+r+\Delta r))^{\frac{1}{\sigma}}}. \tag{A.7}$$

To see that this indeed provides the correct specification of ℓ^I, follow the steps in appendix A.1, using equation (7) instead of equation (A.3). Since the marginal buncher is indifferent between ℓ^τ and ℓ^I it holds that $\mathcal{U}^N - \mathcal{U}^I = 0$, such that

$$\begin{aligned} &\frac{1}{1-\sigma}\left(W_0-(1-\ell^\tau)p\right)^{1-\sigma} + \delta\frac{1}{1-\sigma}\left(p[1-\ell^\tau(1+r)]\right)^{1-\sigma} \\ &-\frac{1}{1-\sigma}\left(W_0-\left(1-\frac{1+\left(1-\frac{W_0}{p}\right)(\delta(1+r+\Delta r))^{\frac{1}{\sigma}}}{(1+r+\Delta r)+(\delta(1+r+\Delta r))^{\frac{1}{\sigma}}}\right)p\right)^{1-\sigma} \\ &-\delta\frac{1}{1-\sigma}\left(p\left(1-\frac{1+\left(1-\frac{W_0}{p}\right)(\delta(1+r+\Delta r))^{\frac{1}{\sigma}}}{(1+r+\Delta r)+(\delta(1+r+\Delta r))^{\frac{1}{\sigma}}}(1+r+\Delta r)\right)\right)^{1-\sigma} = 0. \end{aligned} \tag{A.8}$$

This implicit equation implicitly defines the elasticity of mortgage demand. However, the the values of W_0, p and δ are unknown. Hence, it is not feasible to retrieve the elasticity of mortgage demand via this route.

Appendix B: Supplementary figures and tables

Figure B.1: LTV distributions ABN AMRO & SNS

Panel A: *The LTV distribution for ABN AMRO*

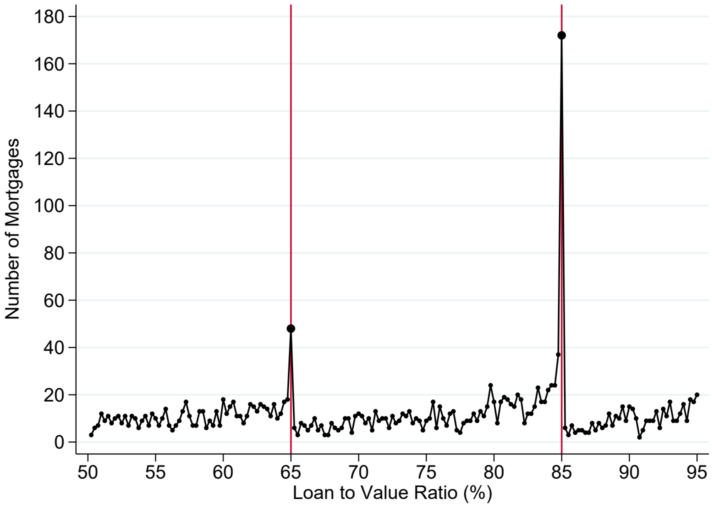

Panel B: *The LTV distribution for SNS*

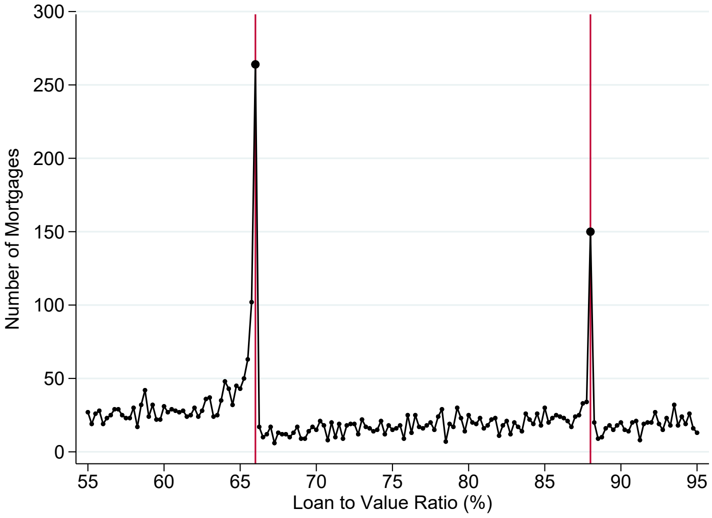

Notes: These graphs display the observed distribution of the LTVs of Dutch households that have taken out a mortgage at ABN AMRO (or an affiliated label) and at SNS (or an affiliated label). For ABN AMRO, interest rate notches are located at the LTV thresholds 65 and 85. The SNS notches are at 66 and 88.

ABN·AMRO < ABNAMRO.nl

vrijdag 20 september 2019

Hypotheekproduct: Budget Hypotheek
Hypotheekvorm: Annuïteiten
Tariefklasse: t/m 85%

- [x] Huisbankkorting 0,2%
- [ ] Duurzaamheidskorting 0,2%

Looptijd	NHG	t/m 65%	t/m 85%	meer dan 85%
Variabel	1,45 %	1,50 %	**1,65 %**	1,90 %
1 jaar	1,24 %	1,32 %	**1,47 %**	1,72 %
2 jaar *	1,24 %	1,32 %	**1,47 %**	1,72 %
3 jaar *	1,29 %	1,37 %	**1,52 %**	1,77 %
5 jaar	1,32 %	1,43 %	**1,58 %**	1,78 %
6 jaar	1,32 %	1,43 %	**1,58 %**	1,78 %
7 jaar *	1,33 %	1,44 %	**1,59 %**	1,79 %
10 jaar	1,34 %	1,48 %	**1,64 %**	1,79 %
12 jaar *	1,62 %	1,67 %	**1,84 %**	2,04 %
15 jaar	1,83 %	1,94 %	**2,04 %**	2,27 %
17 jaar *	1,88 %	1,99 %	**2,09 %**	2,28 %
20 jaar	1,90 %	2,04 %	**2,14 %**	2,31 %
30 jaar	2,30 %	2,40 %	**2,50 %**	2,75 %

Figure B.2: ABN AMRO's reported mortgage interest rate

Notes: Actual (20-09-2019) mortgage interest rates reported by ABN AMRO.

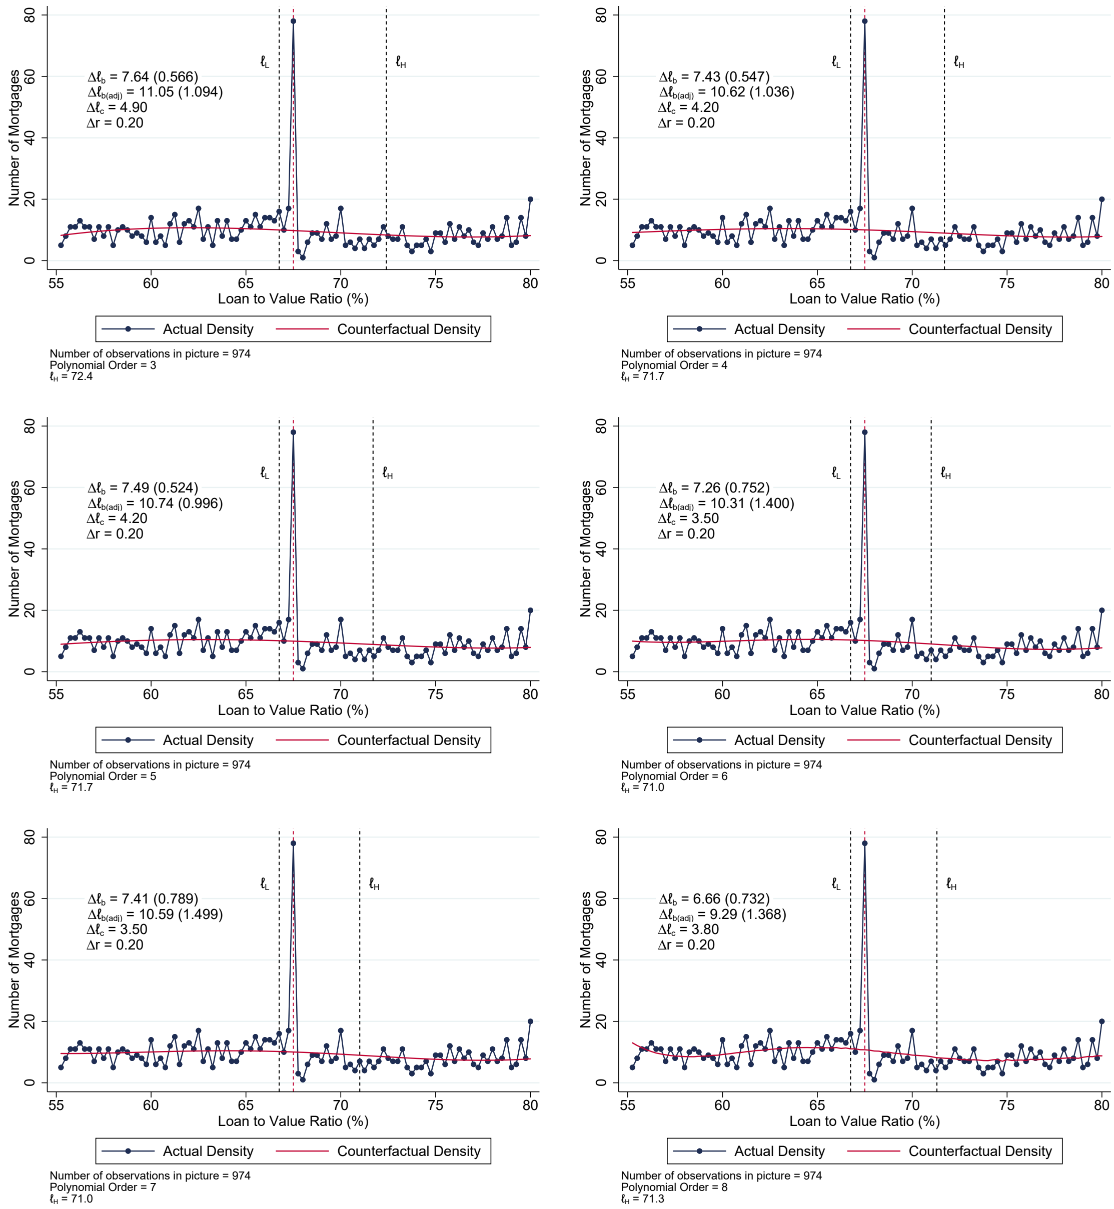

Figure B.3: Actual and counterfactual LTV distribution around the 67.5 notch for Rabobank

Notes: See for an additional description the notes under Figure 5.

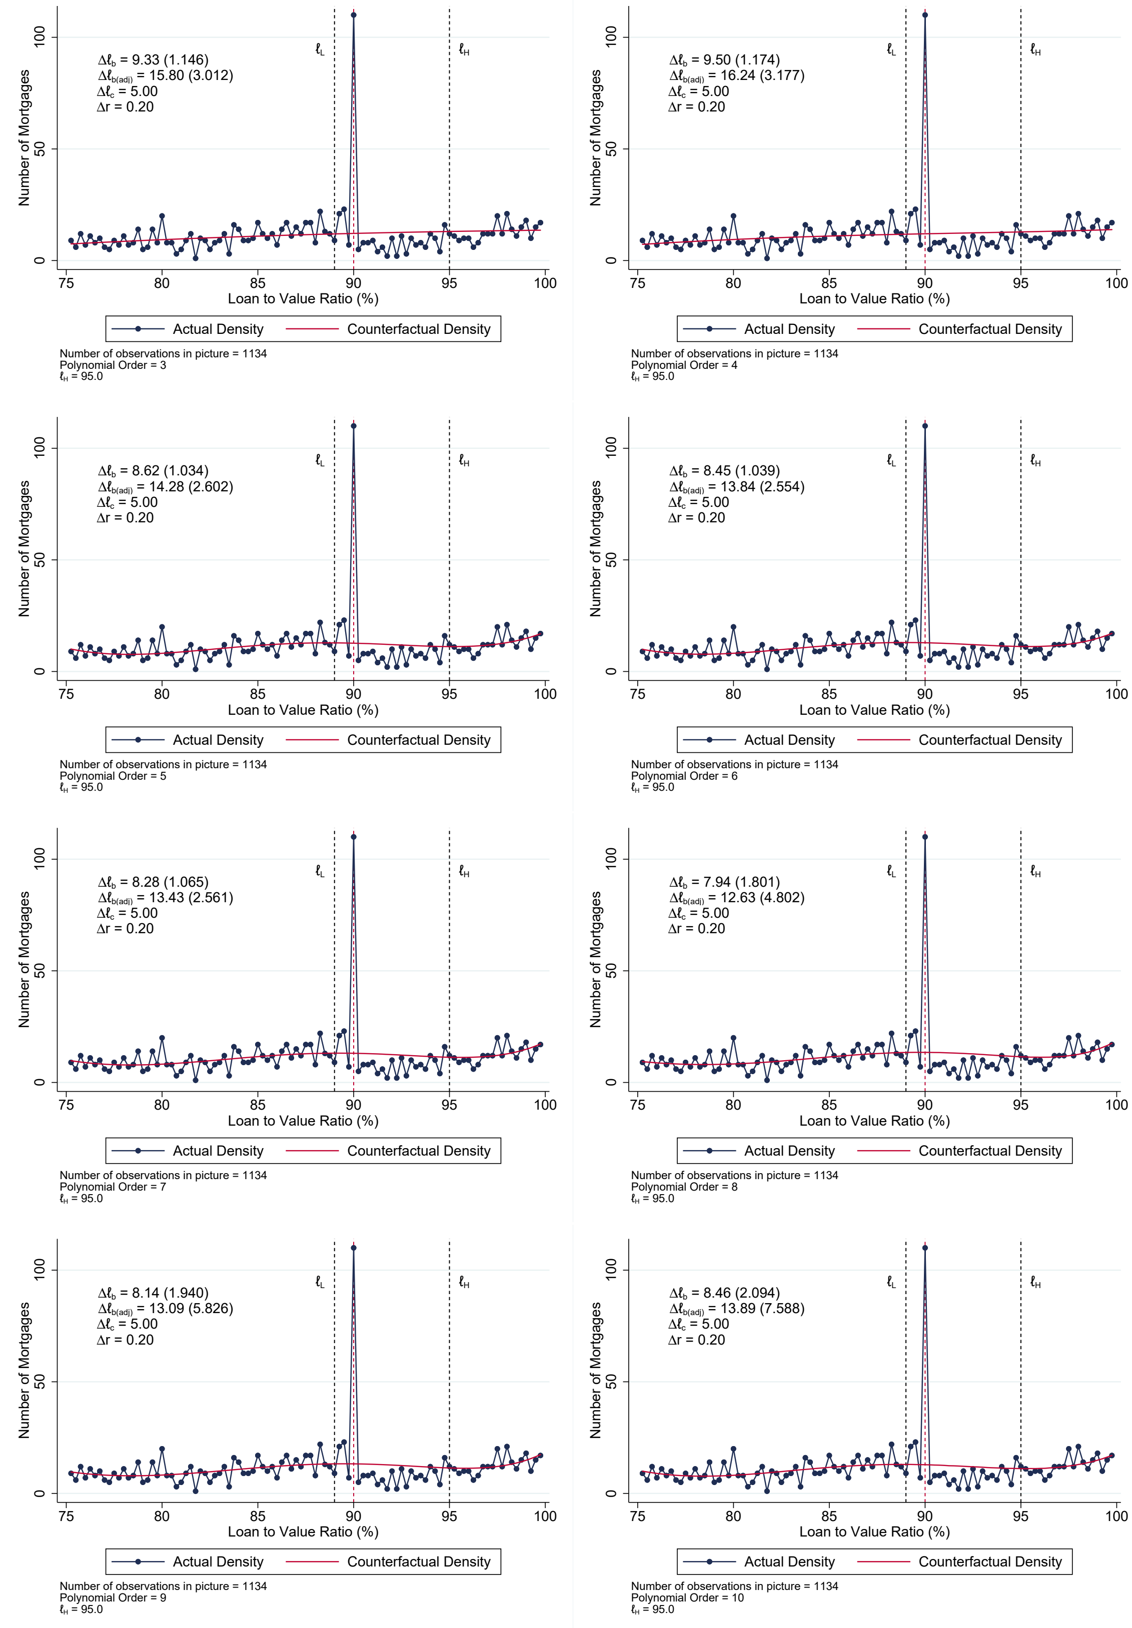

Figure B.4: Actual and counterfactual LTV distribution around the 90 notch for Rabobank

Notes: See for an additional description the notes under Figure 5.

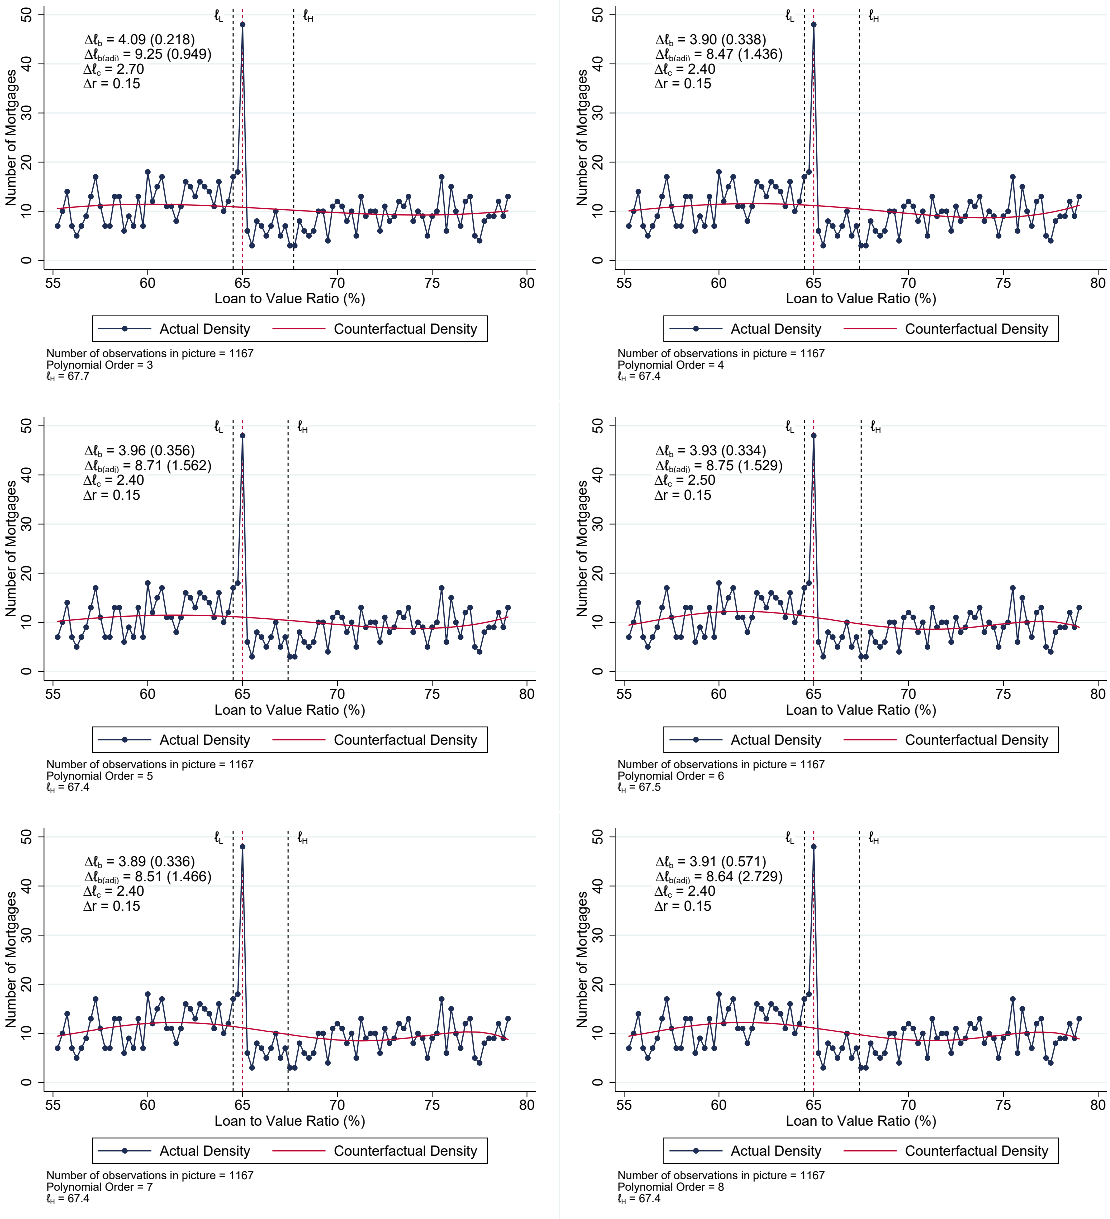

Figure B.5: Actual and counterfactual LTV distribution around the 65 notch for ABN AMRO

Notes: See for an additional description the notes under Figure 5.

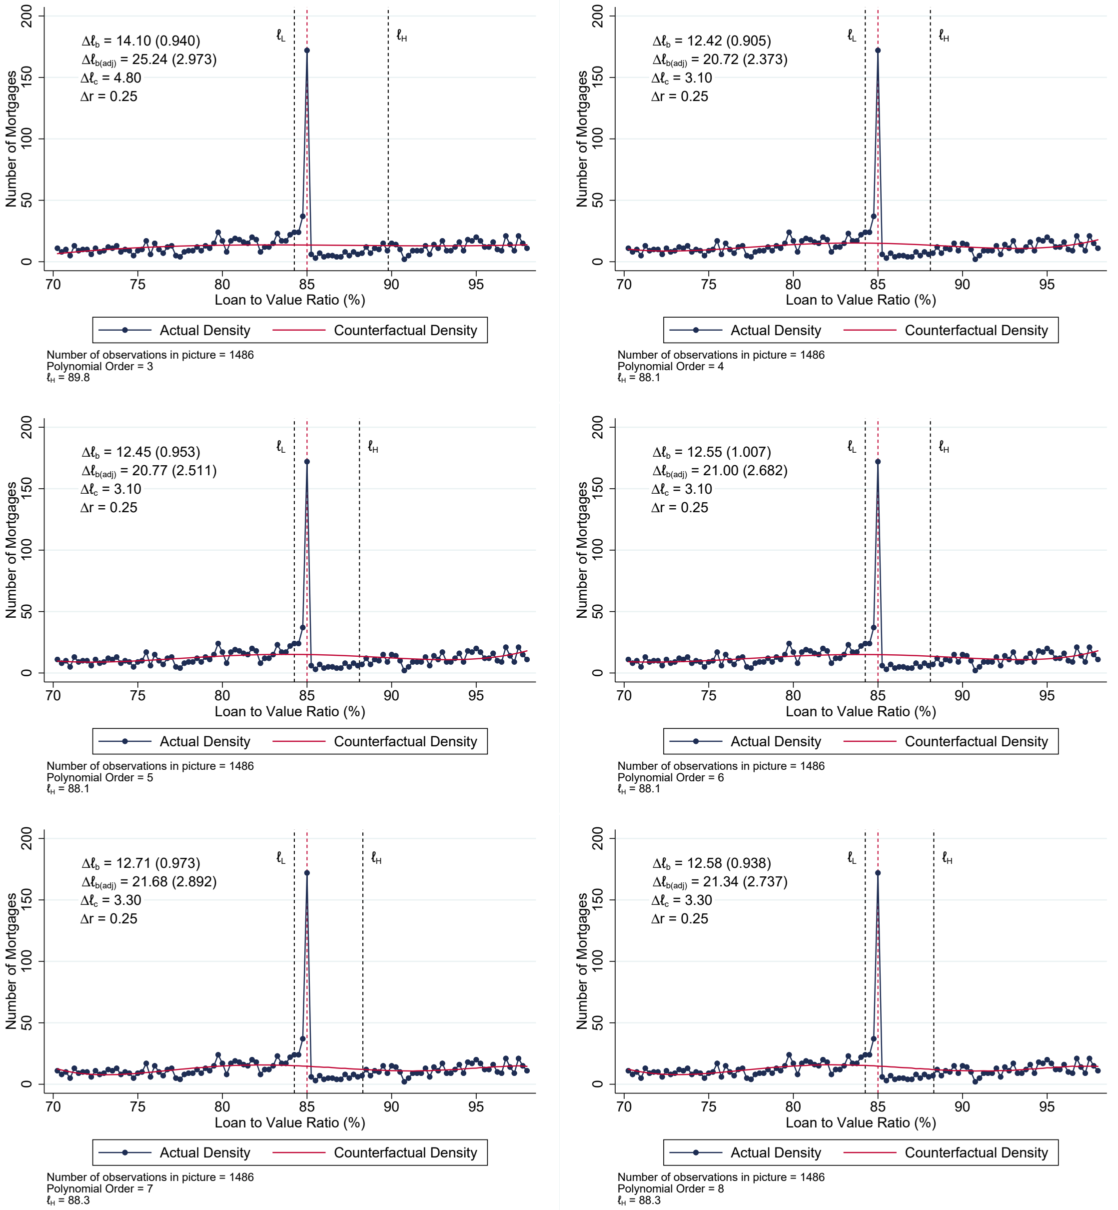

Figure B.6: Actual and counterfactual LTV distribution around the 85 notch for ABN AMRO

Notes: See for an additional description the notes under Figure 5.

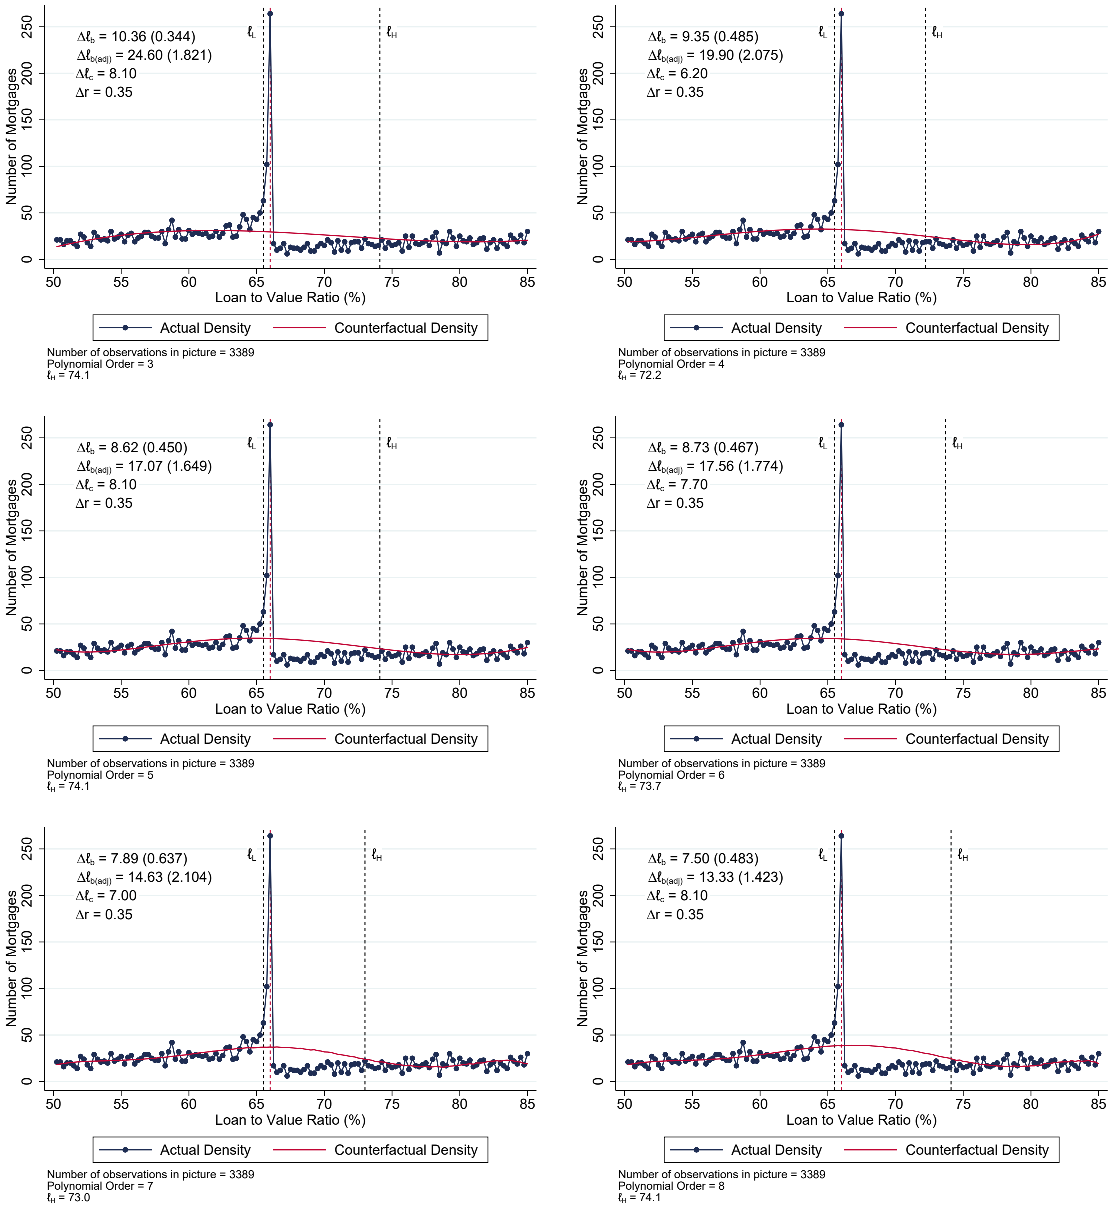

Figure B.7: Actual and counterfactual LTV distribution around the 66 notch for SNS

Notes: See for an additional description the notes under Figure 5.

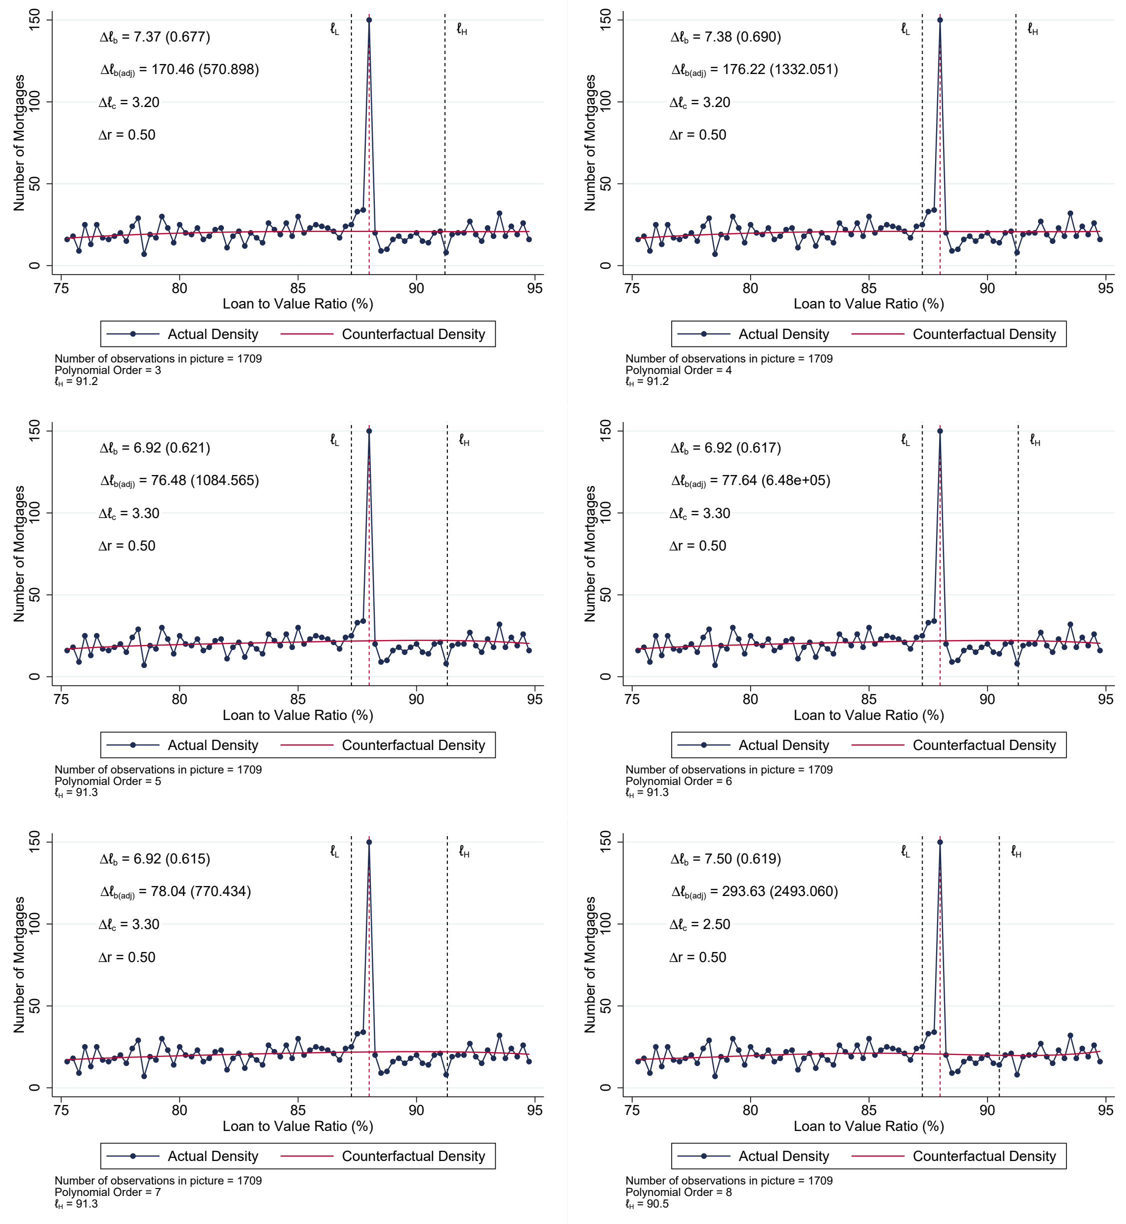

Figure B.8: Actual and counterfactual LTV distribution around the 88 notch for SNS

Notes: See for an additional description the notes under Figure 5.

Appendix C: Glossary

DNB De Nederlandsche Bank

EIS elasticity of intertemporal substitution

LLD Loan Level Data

LTV loan-to-value ratio

MID mortgage interest deduction

RDD regression discontinuity design

RKD regression kink design