BANK OF ENGLAND

Staff Working Paper No. 808

Modelling the distribution of mortgage debt

Iren Levina, Robert Sturrock, Alexandra Varadi and Gavin Wallis

July 2019

BANK OF ENGLAND

Staff Working Paper No. 808

Modelling the distribution of mortgage debt

Iren Levina,[1] Robert Sturrock,[2] Alexandra Varadi[3] and Gavin Wallis[4]

Abstract

This paper presents an approach to modelling the flow and the stock of mortgage debt, using loan-level data. Our approach allows us to consider different macroeconomic scenarios for the housing market, lenders' and borrowers' behaviour, and different calibrations of macroprudential policy interventions in a consistent way. This, in turn, allows us to take a forward-looking view about potential risks stemming from the distribution of mortgage debt, as well as assess the impact of potential macroprudential policies in a forward-looking manner.

Key words: Mortgage market, housing market, macroprudential policy, loan-level data, flow model, stock model.

JEL classification: D04, G21, R20, R21, R31.

(1) Bank of England. Email: Iren.Levina@bankofengland.co.uk
(2) Formerly Bank of England
(3) Wadham College, Oxford
(4) Formerly Bank of England

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. The models and associated outputs described in this working paper reflect the work of a number of people over the past few years including Katie Low, Sagar Shah, Vas Madouros, Tom Stratton and Mariana Gimpelewicz. We are also grateful to an external referee for comments and feedback.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

For most individuals a mortgage is the largest financial transaction they will ever make, one that often requires them to become significantly leveraged. Given the role played by highly indebted households in financial crises (Jorda et al 2016, Mian and Sufi 2010), it is unsurprising that central banks focus intently on the mortgage market.

Despite widespread house price falls during the global financial crisis, in recent years many countries have seen a renewed growth in house prices. The IMF's Global House Price Index, an average of real house prices across 57 countries, has almost returned to its pre-crisis level. In the UK, the house price to income ratio has recently reached its 2006 average. This has fuelled a rise in the share of new mortgages with high loan to income multiples over time.

The continued growth in housing activity and sustained increases in house prices relative to incomes have prompted policymakers to look more closely at the distribution of debt rather than focusing on simple averages to measure vulnerabilities in the household sector. That is because average or aggregate measures can miss important risks arising from the tail of the household debt distribution, which can amplify economic shocks. Highly indebted households – i.e. those in the tail of the household debt distribution – are particularly vulnerable to unexpected events that increase the burden of servicing existing debts, such as an increase in interest rates or a fall in incomes. For example, in the UK more heavily indebted households cut spending by more during the 2007-09 crisis than did less indebted households (Bank of England 2017). And although the majority of households take mortgages with loan to income ratios below 4, the share of new mortgages above 4 has been rising consistently since the crisis and now exceeds its peak in 2007.

To assess risks stemming from household indebtedness, it is important not only to have a picture of the current distribution of debt, but also to understand how it may evolve under certain economic conditions. For example, the ability of some households to service their debt could be challenged by a period of weaker employment and income growth. These vulnerable households could affect broader economic activity by cutting back sharply on expenditure in order to continue to service debts. Alternatively, vulnerable households may default on their debts, testing the resilience of lenders. Being able to model the possible distribution of household debt under various scenarios is of paramount importance for policymakers when trying to quantify the associated financial stability risks to the economy.

Understanding how risks from the distribution of household debt may evolve over time is also important for policy development. Since the global financial crisis many countries have introduced macroprudential household sector tools (see, e.g., Jácome and Mitra 2015). Among these tools, restrictions on loan to value (LTV) and debt-service-to-income ratios are the most frequently used globally (IMF 2018). Some countries also have restrictions on high-loan to income (LTI) lending (e.g., Australia and Ireland). In the UK, in June 2014 the Bank of England's Financial Policy Committee (FPC) announced a LTI flow limit, aimed at insuring against the risk of a marked loosening in underwriting standards and a further significant rise in the number of highly indebted households. Under that policy, mortgages with an LTI ratio at or above 4.5 may not exceed 15% of lenders' new mortgage lending. Lenders were also required to ensure that borrowers were able to afford their mortgage even if Bank Rate rose by 300 basis points (bps) during the first five years of the mortgage.[1]

Assessing the impact of these policies is difficult. It is relatively simple to establish which recently extended loans would not have been allowed under a given macroprudential policy. For example, if a central bank or government requires no loan be extended at above a 95% LTV ratio, these loans will cease to exist. But more difficult questions remain. How will banks and borrowers respond to these restrictions? How will the impact of these limits change over time and under different macroeconomic conditions? Without answers to these questions it is impossible to estimate the full effect of macroprudential policies. Yet, an ability to arrive at such estimates is important. They could affect the calibration of the policies, if policy makers want to be forward looking in their risk assessment and if they want to set policy in a way that accounts for potential future changes in the distribution of mortgage debt.

This paper outlines our approach to producing a loan-level projection of UK mortgage debt. This projection allows us to estimate the risks from household indebtedness, as well as to assess the impact of macroprudential policies, under a given macroeconomic scenario, policy calibration, and assumptions about lenders' and borrowers' behaviour. By using loan-level data on the entire UK population of owner-occupier mortgage loans we are able to construct

[1] In June 2017 the FPC clarified its affordability test Recommendation, to promote consistency in its implementation across lenders. The Committee made clear that lenders should apply a 300bps buffer over the reversion rate, i.e. the rate the borrowers would revert to at the end of their fixed term contract, unless they switch to a new contract. For most lenders, it's the standard variable rate (SVR).

distributions of the key mortgage characteristics rather than look at simple macro aggregates.

There are several steps in our modelling, summarised in Figure 1. First, we project new mortgage lending conditional on a set of macroeconomic variables (a macroeconomic scenario), underwriting standards, borrower behavioural rules, and macroprudential policies. This projection is then combined with an estimate of the current mortgage stock, and both projected new loans and loans within the current stock are updated to reflect changes that occur during the scenario horizon. Fluctuations in borrowers' incomes, principal repayments, house price appreciation, and remortgaging are all accounted for. This produces a loan-level estimate of the future mortgage stock under a given scenario and a set of assumptions. Crucially, the model allows us to vary a macroeconomic scenario and the behaviour of lenders, borrowers and the central bank. This flexibility helps us assess potential risks from the mortgage market and policy impacts in different states of the world.

Projecting the mortgage market at a loan-level is possible because of extensive micro data on the UK mortgage market. The Financial Conduct Authority's (FCA) loan-level Product Sales Data (PSD), available for both new transactions and the current stock of lending, provides the backbone for our modelling. We also draw on the British Household Panel Survey (BHPS), using quantile regression techniques to estimate the impact of macroeconomic conditions on the distribution of borrowers' income changes over time. This element is necessary to properly estimate the number of households vulnerable to unemployment shocks and repayment issues.

We use two illustrative scenarios to demonstrate the model's functionality. One of them is a hypothetical central scenario. The other is an upside scenario, in which the housing market experiences a boom, with house prices and approvals rising more rapidly than expected in the central case. We demonstrate the potential impact of macroprudential policy on the state of the mortgage market in these two hypothetical scenarios, and examine the implications for loan to income (LTI), debt-service (DSR), and loan to value (LTV) ratios.

This paper makes a novel contribution to the literature. Some studies have modelled the distribution of mortgage debt under different macroeconomic scenarios (e.g., Zabai 2017). Other studies have estimated the impact of macroprudential housing tools. For example, Lozej and Rannenberg (2018) have used a DSGE model to investigate the impact of the housing tools on household leverage in Ireland. And the Reserve Bank of New Zealand (2017) has estimated the impact of restrictions on debt to income ratio on the distribution of mortgage debt. To the best of our knowledge, this paper is the first to provide a comprehensive model for projecting the mortgage market at a loan-level in such a way that allows to consider different macroeconomic scenarios, lenders' and borrowers' behaviour, and different calibrations of macroprudential policy interventions in a consistent way.. As such, the model could be invaluable for policy assessment, calibration, and cost-benefit analysis, though these are not the focus of this paper.

The paper is structured as follows. Section 2 outlines our method for projecting new mortgage lending. Section 3 describes how this projection is combined with our estimate of the current stock, updated, and appropriately weighted to produce an estimate of the future stock conditional on a scenario. The data underpinning each of these stages is discussed within the relevant section. Section 4 presents examples of the model's distributional projections under different scenarios. Section 5 concludes.

Figure 1. Model Road Map

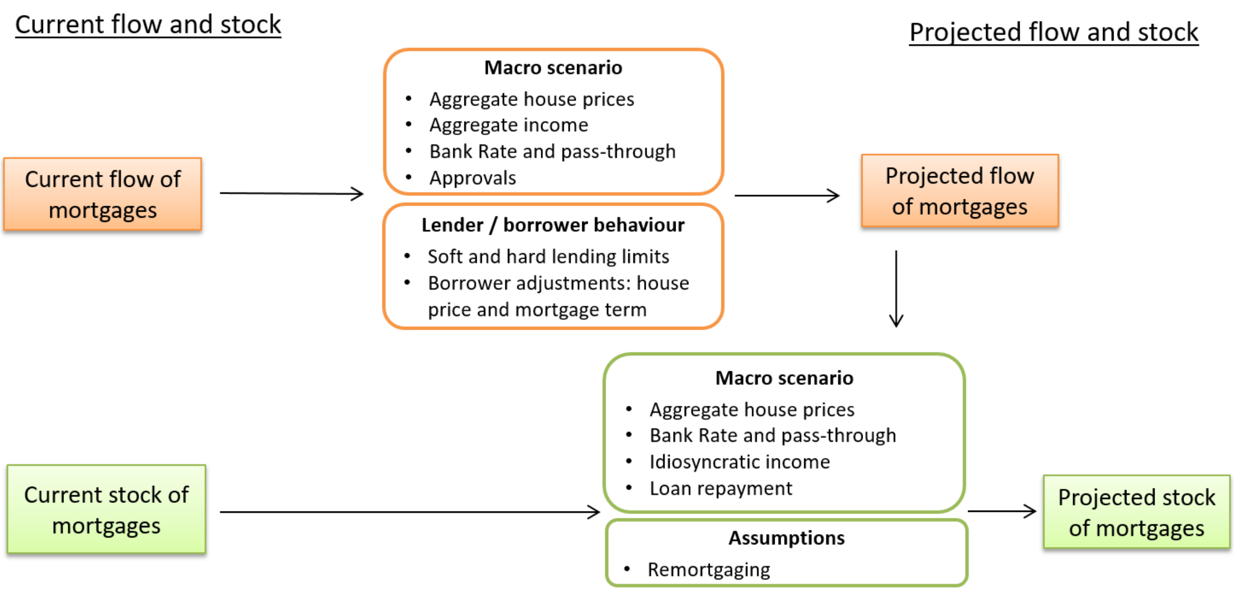

2. Projecting the flow of new mortgages

New mortgage transactions play an important role in the evolution of the mortgage stock and its resilience to shocks. This section describes how we forecast future mortgage transactions using the mortgage flow projection model (the flow model, hereafter).

2.1. Data

The flow model uses mortgage transactions data from the FCA Product Sales Database (PSD). The PSD includes data on new mortgages issued by all financial institutions – banks, building societies, and non-bank finance companies – on a quarterly basis. The database includes important information about mortgage contracts, including loan size, the value of the property purchased, the interest rate, the length of any interest rate deal, and the maturity of the loan. The PSD also captures borrower characteristics such as income, age, employment type (employed, self-employed, retired, or other), and whether a borrower is a first time buyer, home mover or remortgagor. If the loan is a remortgage the PSD captures whether extra money was raised and for what purpose. The PSD does not have information about lenders' reversion rates, which we add to the dataset from Moneyfacts. A reversion rate is a rate a borrower would revert to at the end of their fixed term contract, unless they switch to a new contract. It's a standard variable rate (SVR) for most lenders.

Transactions missing critical information (such as loan value or borrower income) or with implausible recorded values are excluded from our estimation procedure. The total number of transactions excluded is small (usually less than 5%), of which the majority are missing values. The observations missing information are not skewed towards a particular group (e.g., high value properties).

2.2. Methodology: From macro to micro

In order to project future transactions, the flow model takes the most recent four quarters of mortgage transactions from the PSD and adjusts their characteristics in line with a macroeconomic forecast for the following variables:

a. aggregate household disposable income growth;
b. aggregate house price growth;

c. approvals for first time buyers, home movers, and remortgages[2];

d. Bank Rate.

The relationship between the macroeconomic forecast and changes in the projected loan-level characteristics is mechanical. Individual property prices and borrowers' incomes are projected by scaling their current values in line with the macro forecast. Individual mortgage rates and reversion rates are projected in line with changes in Bank Rate, assuming a given pass-through.[3] In the case of households' deposits, we assume that first time buyers save for a deposit out of income, and home movers and remortgagors out of their equity stake in their current property, supported by house price appreciation. We also assume that over the forecast horizon these relationships mirror recent trends. Specifically, we estimate the average relationship between the growth rate in borrowers' deposits and income (for first time buyers) and deposits and house prices (for home movers and remortgagors) between 2012 and 2016. We assume that this relationship would hold over the forecast horizon, although in principle we could vary this assumption depending on a scenario.

Each projected quarter is modelled by taking a representative sample of loans over the four most recent quarters, and adjusting their characteristics in line with a macro scenario, as described above. To illustrate, if an original transaction from 2016 Q1 is being used to project a transaction in 2018 Q1, then original house price in 2016 Q1 will be scaled by the projected growth in aggregate house prices between 2016 Q1 and 2018 Q1:

$$pv_{i,2018Q1} = pv_{i,2016Q1} * \left(\frac{HP_{2018Q1}}{HP_{2016Q1}}\right)$$

where $pv_{i,t}$ denotes the property value for transaction i in period t, and HP_t is the UK aggregate house price index in time period t. These adjustments are illustrated in Figure 2.

These adjustments collectively create a pool of potential new mortgages for each projected quarter. Using this pool, for each quarter and borrower type (i.e. first time buyer, home mover or remortgagor), mortgages are drawn at random, with replacement, to meet the approvals profile in the macro scenario. This step effectively creates a draft set of transactions – the ones that have not yet been evaluated by lenders to see if they are viable mortgages.

[2] The Bank does not forecast the number of remortgages, only their total value. To derive the number of remortgages, recent remortgage approvals are assumed to grow in line with remortgage value in a particular quarter, and then rounded to the nearest whole number.

[3] The model allows us to vary the pass-through.

Figure 2. Illustrative adjustment in line with macro scenario

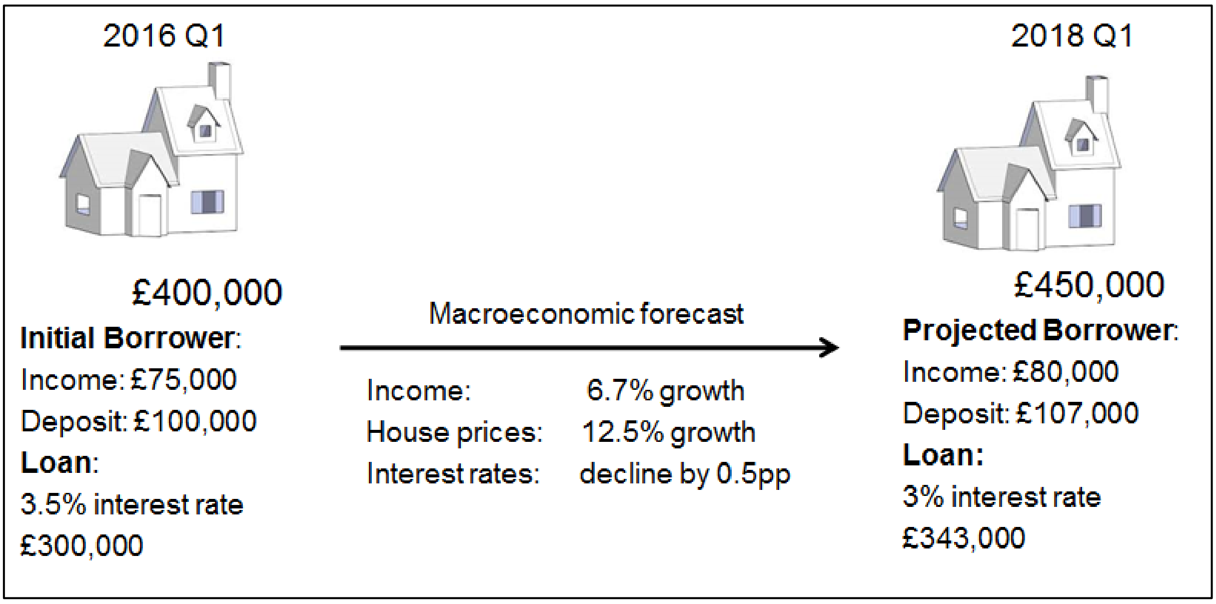

2.3. Methodology: Quality Adjustments

At this stage it is also possible to make more direct adjustments to the underlying quality of borrowers. For example, to reduce the quality of applicants a reduction in income can be applied to some fraction of borrowers in the pool, while holding their remaining characteristics fixed.

This adjustment to certain borrowers partially offsets a key modelling gap. Because the PSD only includes successful mortgage applications, future projected transactions are adjusted versions of transactions from this successful group. In a scenario where underwriting standards remain constant or tighten this is a reasonable modelling approach, as the same borrowers rejected today will still be rejected in the future. However, if underwriting standards loosen, this approach is likely to underestimate the speed or extent to which the quality of the flow of new mortgages deteriorates.

By creating a pool of lower quality borrowers – who are rejected under current standards – we could allow for a loosening in underwriting standards to have a more direct impact on the projected flow. We have previously used this feature of the model for analysis, although it is not used this time to create the illustrative outputs shown in Section 4 below.

2.4. Behavioural assumptions: Borrower and lender behaviour

The second modelling stage involves iteratively testing the draft profile of loans against banks' lending standards. When a loan does not meet the lending criteria, borrowers are allowed to make adjustments. The model allows us to vary the lender policies and borrower behavioural adjustments, as set out below.

Lender behaviour

Lenders are assumed to have lending standards independent of any central bank policy. These lending standards are set as restrictions on loan characteristics, modelled either as a flow limit or a cap. We assume that banks could target such loan characteristics as LTV, LTI, and DSR[4], defined as:

$$LTI = \frac{lv}{y}$$

$$DSR = LTI * \left(\frac{i}{1-(1+i)^{-t}}\right)$$

$$LTV = \frac{lv}{pv}$$

where lv is loan value, pv is property value, y is gross income, i is the relevant interest rate (product or stressed for DSR and stressed DSR, respectively) and t is the mortgage term.

These lending standards allow for a wide variety of potential underwriting criteria. Banks could cap mortgage LTIs at 5, implying any borrower with an LTI above 5 must either adjust their loan or be denied a mortgage. Alternatively banks might set an internal LTI flow limit, for example, allowing only 5% of new loans to have an LTI above 4.75. In this case, if the draft profile of loans exceeds this internal LTI flow limit, some borrowers will be asked to adjust until the final profile of loans does not breach the limit. The model allows us to set some combination of underwriting criteria across LTI, DSR, and LTV.

For simplicity, all loans are tested against a single set of lending standards, rather than modelling separate lending standards for each bank. Because borrowers can apply to several banks, and will be offered a mortgage even if they only pass the criteria of the lender with

[4] DSR can be understood as monthly mortgage repayment as a share of monthly income. For example a borrower with a DSR of 30% spends 30% of their gross monthly income on mortgage repayments. In practice lenders use more sophisticated affordability testing than this simple ratio.

the loosest standards, this method should approximate market outcomes while reducing complexity.

Borrower behaviour

Borrowers are allowed to make adjustments when they fail to meet banks' lending standards. The model permits two possible adjustments. First, borrowers may search for more affordable housing, modelled here as finding a house with a price up to 10% lower. Settling on a cheaper property allows constrained borrowers to reduce their loan size. With a smaller loan, holding income and other characteristics fixed, the borrower is more likely to meet lending standards. However, we assume that only a fraction of borrowers is willing and able to find more affordable housing that meets the lending criteria. We usually use a neutral assumption and set this fraction to 50%.

For a constrained borrower who finds cheaper housing the following adjustments occur:

$$pv_{new} = 0.9pv_{old}$$

$$lv_{new} = lv_{old} - 0.1pv_{old}$$

$$LTI_{new} = \frac{lv_{new}}{y}$$

$$DSR_{new} = LTI_{new} * \frac{i}{1-(1+i)^{-t}}$$

$$LTV_{new} = \frac{lv_{new}}{pv_{new}}$$

The second possible adjustment allows borrowers bound by the DSR limit to increase their mortgage term up to a maximum of 35 years, provided doing so would not result in their loan extending into retirement, modelled here as a maximum age of 68. While both mortgages at terms above 35 years and lending into retirement do occur, they are uncommon.

If a borrower still fails any of the lending standards after these adjustments they are dropped from the approvals profile and replaced by a new borrower who passes. This process continues until nobody breaches lenders' underwriting standards and the number of loans approved matches the approvals profile in the macro scenario, consistent with the assumption that our forecast scenario is conditioned on lenders' underwriting standards within that scenario.

Table 1 summarises the calibration of the lender and borrower behaviour for our two illustrative scenarios. We assume that lenders have internal LTI and DSR limits. For the purposes of this example, we do not impose any constraints on LTV, although in principle our modelling approach could allow for that as well.

Table 1. Banks' lending standards and borrowers' adjustments under different scenarios

Behavioural assumption	Central	Upside
LTI limit	5	6
DSR limit	41	50
Interest rate used by lenders for testing a borrower's affordability	6.2	5.5
Max reduction in house price when a borrower is above LTI or DSR limit	10%	10%
Max mortgage term to which borrower can extend when above DSR limit	35	35

2.5. Incorporating macroprudential policy

Macroprudential policies are applied to the profile of loans created after banks' lending standards and borrowers' adjustments have been accounted for. Similar to banks' lending standards, the macroprudential policy is modelled as caps or flow limits on specific loan characteristics, for example, an LTI flow limit or a stress rate used in affordability assessment. Borrowers that fail these policy limits are allowed to adjust their loans in order to pass, unless they already adjusted to meet banks' internal standards. If, after the adjustments, the borrowers still breach the macroprudential policy, their mortgage application is assumed to be rejected and the loan is dropped out of the modelling sample.

In practice, both banks' and macroprudential standards would be applied simultaneously by a bank. But applying them separately in the model is necessary because it allows us to isolate the impact of lenders' internal standards and macroprudential policy, and hence to estimate the macroprudential policy impact on the number of future mortgage transactions and their characteristics.

3. Projecting the stock of mortgages

While creating a projection of the flow of new mortgages provides useful information, it is only a part of the picture. To fully understand the direction of the mortgage market and assess changing risks and vulnerabilities it is also necessary to understand what is happening to existing mortgages and borrowers. Collectively, this stock of mortgages provides a more complete picture of how debt and affordability are expected to change in aggregate, which is more useful for policy. It is not just the response of new borrowers to macroeconomic and financial shocks that matters, rather the impact of those shocks on the entire mortgage market, and the corresponding spill-overs to the wider economy.

Conceptually, there are two essential features of projecting forward a stock of mortgages. First, a projection should be stock-flow consistent, i.e. the projected stock should be a sum of the current stock and projected flow. Second, a projection should take into account changes in the loan characteristics and macro scenario over time.

In practice, these two considerations imply a three-step modelling process: estimating the existing stock, projecting forward a flow of new mortgages, and adjusting both in line with the macro scenario. Estimating the existing stock of mortgages relies on the FCA's Product Sales Database – 007 Mortgage Performance data (henceforth PSD007), but due to the limitations of this dataset a matching and cleaning procedure must be used in order to make it suitable for our modelling. This process is described below. The second step, projecting forward a flow, has been covered in detail in Section 2 above. Finally, these two elements are combined and adjusted to create a simulation of the future UK mortgage stock. Adjustment is necessary because loan and borrower characteristics may change substantially from origination, or in the case of the stock data, from our current snapshot. Most borrowers will have partially repaid their mortgages, reducing their outstanding balance. House prices may have risen or fallen, changing borrowers' leverage, and changes in interest rates may result in altered payments for borrowers upon remortgaging. Finally, borrowers' incomes may change, as a result of promotion or being laid off, altering their ability to make mortgage payments. All of these possibilities must be accounted for and linked to a macroeconomic scenario.

The data and methodology used to create the existing mortgage stock is outlined briefly in Section 3.1 and uses previous work.[5] Sections 3.2-3.3 outline how we combine and adjust the existing stock and projected flow of new mortgages to create a projection of the future mortgage stock.

3.1. Estimating the current stock

Data on the existing UK mortgage stock (up to 2015 Q2 at the time of estimation) is available through the PSD007. This data covers all outstanding regulated UK mortgages and contains information on several loan characteristics, including the outstanding mortgage balance, current rate of interest, interest rate type (fixed or floating), and mortgage performance (arrears, default). It also includes detailed identification information on the loan, including the property post code, the date the mortgage was originated, and the original loan size.

However, the PSD007 is missing several key variables required to calculate borrower leverage and affordability measures. It does not contain the property value (either at origination or a current assessment), the type of loan (e.g. first time buyer, home mover, remortgage), repayment type (e.g. capital and interest, interest only) or borrower income (either at origination or a current assessment).

These characteristics are included in the FCA's Product Sales Database – 001 Mortgage (henceforth PSD001), which contains detailed mortgage information (as described in Section 2.1) for all regulated UK mortgage transactions from 2005 Q2. Because most of the outstanding UK mortgages in 2015 Q2 will have either been originated or remortgaged between 2005 Q2 and 2015 Q2, it should be possible to find a matched entry in PSD001 for most mortgages in PSD007. The datasets are matched using several loan characteristics present in both datasets, including post code, original loan value, borrower's age, and the date the mortgage was opened. As a result it is possible to match more than 75% of the mortgages in PSD007. This matching process is discussed in detail in Chakraborty et al (2017), while section 3.1.1 briefly summarises the authors' approach to creating an estimate of the 2015 Q2 UK mortgage stock.

The matched stock closely matches measures of the UK mortgage stock available from survey evidence[6], and there is no consistent feature that distinguishes matched observations

---

[5] See Chakraborty et al (2017).
[6] See Chakraborty et al (2017, p. 11-14).

from those that remain unmatched. As a result, we are confident that the matched stock dataset provides an unbiased sample of the true UK mortgage stock.

3.1.1. Estimating the stock as of 2015 Q2

At the time of estimation, the latest loans available in PSD007 were originated in 2015 Q2, but most are legacy loans originated in previous years. To obtain a consistent and comparable dataset, house prices and incomes are updated for each individual loan, from origination to 2015 Q2.

House prices are updated using the median value of the house price growth in the Land Registry's UK wide house price index, regional indices and a more granular Price Paid dataset between the date of origination and 2015 Q2.[7] Income is updated in two steps. First, we evaluate the distribution of income changes by borrower characteristics between 1991 and 2015, based on data from three household-level panel surveys: the British Household Panel Survey (BHPS), Understanding Society (USoc) and Labour Force Survey (LFS). Specifically, mortgage borrowers in the surveys are grouped into seven mutually exclusive and collectively exhaustive groups based on their income growth. A normal distribution is then fitted on income growth for each group and year. Second, PSD007 is then split into the same seven groups and individual incomes are grown by an income change drawn randomly from the survey data distributions for the corresponding group and year (with borrowers being able to switch from one group to another over time).

While there is no PSD back data to validate this income projection, Chakraborty et al (2017) compare the resulting PSD income distribution against the Wealth and Asset Survey (WAS), which they believe provide the best available sample of household income across mortgagors. Figure 3 shows that the two income distributions are broadly comparable.

[7] Although more granular geographical region index can capture regional dynamics, these indices are calculated by us based on sale prices of transactions that take place without property characteristic adjustment. Properties are generally sold when their prices are sufficiently appreciated, so indices based on transactions tend to be biased upwards. Also, smaller geographic areas have fewer properties sold per quarter making the indices more volatile. By contrast, the UK regional index, while less granular, is adjusted for some of these biases. For more detail, see Chakraborty et al (2017, Appendix 6.7).

Figure 3. Income distribution in the matched PSD data and the WAS

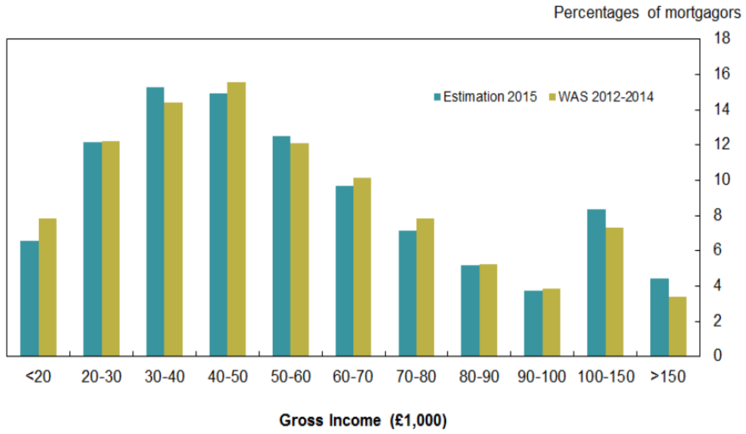

3.1.2. Adding flow data to estimate the current stock

The matched stock data as of 2015 Q2, estimated using the approach described above, does not give us an up-to-date measure of the current stock (end-2016 for the results presented here). To obtain such an estimate we augment the estimated 2015 Q2 stock, with the latest six quarters of flow data available in the PSD001.

3.2. Adjusting stock and flow to future periods

Having estimated the current stock and projected flow, using the methodology summarised above, we need to adjust both before we can produce an accurate measure of the future mortgage stock. These adjustments should reflect changes in borrower and loan characteristics over time (e.g. income, interest rates, outstanding loan balance, and property value), between loan origination (or 2015 Q2 in the case of the stock) and the end of the forecast horizon. These adjustments are applied separately to the 2015 Q2 stock, flow since then, and projected flow, but using the same methodology, summarised in this section.

3.2.1. Modelling income

Projecting forward a mortgage stock, we need to account for changes in borrowers' income over a forecast horizon. These changes should reflect a macro scenario, but it is also important to allow for heterogeneity in these changes. If instead we adjusted individual borrower incomes uniformly in line with aggregate changes to household income, the result would not represent how incomes evolve on a borrower-by-borrower basis. While aggregate nominal income may grow 4% in a given year, the variation for individual borrowers is substantial. Some borrowers' wages will not change, despite increases in inflation. Others will see their take-home pay wage fall more directly, as the result of being moved to part-time work, being laid off, or in the case of self-employed individuals, a year of bad business. Some individuals will be promoted and realise double-digit wage increases. Without this variation few borrowers would ever face repayment difficulties because mortgage payments would always remain near the (affordability tested) share of income they were originated at.

We model the distribution of income changes over a given scenario using a quantile regression, which estimates a distribution of individual income changes, consistent with aggregate changes in income and unemployment. The estimated distribution of individual income changes is then applied to the existing stock of mortgages and our mortgage projections, both of which are created using the PSD data. To do that, we must assume the survey data is a representative sample of PSD. Distributional statistics for both datasets match closely, suggesting this assumption is reasonable (Figure 4). At least some of the observed difference may be due to the fact that the BHPS survey data and PSD007 are from different time periods, 2013 and 2015, respectively.

Figure 4 (a). Distribution by employment status

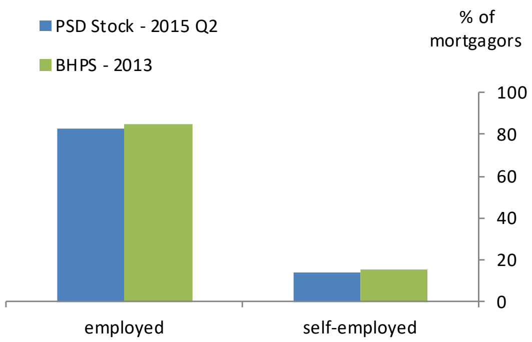

Figure 4 (b). Income distribution

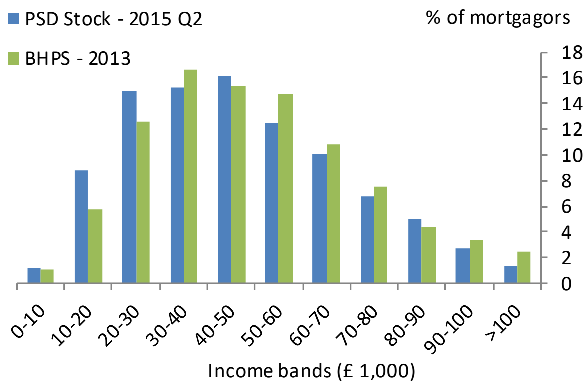

To allow for heterogeneity, we want a distribution of income changes to reflect the borrower characteristics – their income level, whether they are employed or self-employed, and whether they recently experienced an income shock. To capture these differences, we split the individuals in the BHPS survey into 18 cohorts, reflecting a combination of the above characteristics, and project an income change distribution for each cohort. To apply these income changes from the survey sample to the exiting stock and projected flow, we split these loans into the same cohorts as the survey data. Then, for a given income-employment-prior shock cohort, the income changes experienced by borrowers in a given projected year are randomly drawn from the distribution for the same cohort and the same projected year from the estimated survey distribution.

Combining these annual income changes creates a projected income for each individual loan at the end of the forecast horizon:

$$y_{i,projected} = y_{i,initial} * \prod_{t\,=\,initial}^{projected} s_{i,t}$$

where $s_{i,t}$ is the income change for borrower i in time t (where t is a period between origination and the end of the projection). This projection applies to the projected loans and the six quarters of loans originated between 2015 Q2 and 2016 Q4 that were added to the existing stock. For loans in the stock as of 2015 Q2, where income has already been updated to 2015 Q2, the formula is simply a specific case of the one above, where $initial$ is set to 2015 Q2 for all loans.

3.2.2. Modelling other mortgage characteristics

Borrowers' income is just one of the characteristics that has to be updated to project loans to the end of the forecast period. In order to create an accurate measure of the future mortgage stock, we also need to adjust other mortgage characteristics, including interest rates, loan terms, property values, and loan amounts. These adjustments are set out below.

Property values

The property value for a projected mortgage is the initial property value grown in line with the forecast for house price inflation:

$$pv_t = pv_0 * hpi_t/hpi_o$$

where pv_0 is the property value at the time of mortgage origination (or 2015Q2 if the mortgage is part of the existing stock in 2015 Q2), hpi_0 is the house price index at mortgage origination (or 2015 Q2 for the stock data), and hpi_t is the house price index in the projected year t, provided by the forecast for house prices within the scenario. It would be desirable to use more granular house price forecasts in this projection, for example, at a regional level. However this has been left for future work.

Loan values

Mortgage loan values are updated to account for repayments using the standard repayment formula below. This repayment formula assumes equal mortgage repayments, for a given interest rate, over the entire duration of the mortgage.

Loan value for a projected mortgage ($l\mathrm{v}_\mathrm{t}$) is:

$$\mathrm{lv_t} = \mathrm{lv_0} * (1 + \mathrm{r_t})^{\mathrm{m}} - \frac{[(1 + \mathrm{r_t})^{\mathrm{m}} - 1]}{\mathrm{r_t}} * \text{repayments}$$

where lv_0 is the loan value at the starting point, $\mathrm{r_t}$ is the quarterly interest rate in the projected year , m is the quarterly term between the starting point and the projection year, and repayments are calculated as:

$$repayments = \ lv_0 * \frac{r_t * (1 + r_t)^T}{[(1 + r_t)^T - 1]}$$

where T is the total mortgage term. In the case of interest-only mortgages, the outstanding loan value remains fixed over the projection. Our approach does not account for mortgage over-payments, early repayments or missed payments, in other words, we assume that borrowers only make contractual repayments. As a result, we underestimate mortgage repayments, and this is one of the elements of our model that we would like to improve in the future.

Interest rates and mortgage term

Interest rates are updated in line with changes in Bank Rate between 2015Q2 and 2016Q4, and after that – in line with a macro scenario. We usually assume a 100% pass-through from Bank Rate to a mortgage rate, although the model allows us to vary the pass-through. We assume that borrowers with fixed term mortgages ending during the scenario period re-fix their mortgage at the new prevailing interest rate in that time period. We assume that these borrowers have a revealed preference for a contract length, i.e. that a borrower with a

2-year fixed mortgage at origination will remortgage every two years through the forecast period.[8] Rates on mortgages other than fixed-rate are assumed to track Bank Rate.

This approach is straightforward for updating the rates in the current and projected flow, as contract length for the fixed rate mortgages is a good quality field in PSD001. But this field has a lot of missing observations in the PSD007 stock data. To address this issue, in cases where the mortgage deal length is unavailable in the 2015Q2 stock, we estimate when the fixed term of the mortgage will end. We use Bank of England data on average interest rates charged on different fixed rate mortgage contracts to predict the most likely product type of each individual loan. For example, 5-year fixed rate contracts are likely to be on a higher rate than 2-year, 3-year or even 4-year fixed rate contracts. So we match the 2015 interest rates in PSD with the interest rate ranges and corresponding contract length in the matching Bank data.

To calculate DSRs, we also need to adjust the remaining mortgage term, which is estimated as decreasing with the time elapsed since the starting position (or since 2015 Q2 in the case of loans in the 2015 Q2 stock). Again this process assumes no over-payments, early repayments or missed payments.

3.3. Merging the stock and the flow (weighting)

Having adjusted the stock and the flow data to the end of the forecast horizon, using the method described in Section 3.2, the final step in creating a mortgage projection is to bring together the adjusted stock and the flow, accounting for remortgaging and data sampling. We need to avoid double counting mortgagors that would have remortgaged during the forecast horizon – they should only appear once in the stock of mortgage debt. So we remove mortgages out of the current stock and the current flow that would have remortgaged, assuming that interest only mortgages will not remortgage over the forecast horizon. And we assume that half of the mortgagors in the projected flow would also remortgage over the forecast horizon, so we remove them from the projected flow as well. This assumption is based on a fact that in the 2015Q2 mortgage stock a median loan originated about three and a half years ago, and we assume this to hold in the future. This suggests that, on average, a loan would remortgage once over a seven year horizon.

[8] This implies that there is no switch from fixed term contracts to SVRs in the remaining stock over the course of the scenario. We can vary this assumption to allow for some borrowers to revert to their lenders' SVR.

4. Model outputs: Illustrative scenarios

Once merged, the stock, the current flow, and the projected flow can then be used to create a picture of the future stock of UK mortgage debt. Our projections focus primarily on the distributions of variables that are relevant for assessing potential vulnerabilities from household debt – mortgagors' LTIs, LTVs, and DSRs – but because the projection produces a loan-level dataset, in principle it is possible to cut this projected stock along lots of different dimensions.

To demonstrate the model's functionality, we present two hypothetical scenarios – an upside and a central case. We use the flow of mortgage lending in 2016 and consider how the distribution of debt could have evolved over 2017-2023. We assume that in the central scenario the macroeconomic environment and conditions in the housing market are benign over this period. House price growth is close to 5% for a few quarters, slightly weaker than in 2016, before stabilising at just below 4% annually. Nominal household income growth is about 3% in the near term, similar to the 2016 outturns, before recovering to a long run growth rate of 3.4%. So the aggregate house price to income ratio rises only slightly over the forecast. Total approvals strengthen somewhat from around 210,000 to 222,000 over the forecast horizon. Under this scenario we would expect small changes in the stock of UK mortgage debt, mostly through repayment effects.

As a result, the FPC's housing policies – which limit the flow of loans at LTI ratios at or above 4.5 to 15% of new lending and set the stressed rate for the mortgage affordability test to 300bps above the reversion rate (the standard variable rate (SVR) for most lenders) – would be unlikely to have any significant impact in the central scenario.

In the upside scenario, conditions in the housing market change significantly. House prices continue to pick up pace in the short run, growing by around 8% annually towards the end of 2017, before returning to a 6% long run growth rate (consistent with average historic UK house price growth). Nominal income growth, however, remains the same as in the central scenario, driving up the house price to income ratio. And mortgage approvals increase substantially, growing to over 270,000 per quarter by the end of the scenario. Bank Rate remains the same as in the central case.

As a result of the substantial increase in house prices, not accompanied by a commensurate rise in income, we would expect the upside scenario to result in a substantial increase in the number of high LTI loans, a pattern seen during previous UK housing cycles (Figure 5).

Figure 5. Share of new mortgage lending, by LTI bucket

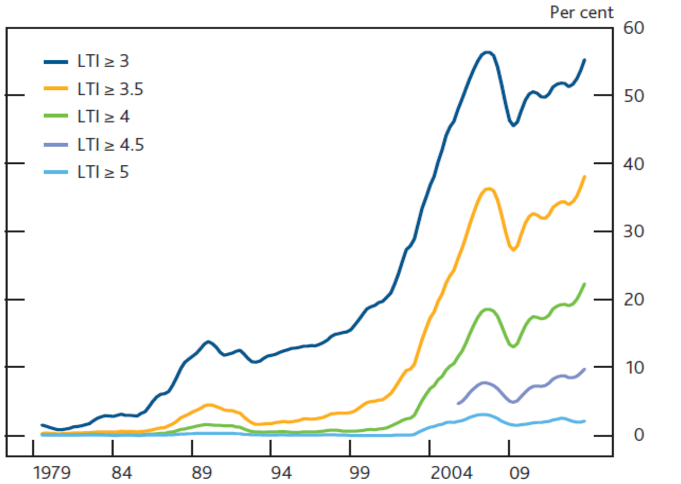

Sources: Council of Mortgage Lenders, FCA Product Sales Data (PSD) and Bank calculations.

(a) See footnotes to Chart 2.9.
(b) Prior to 2005 Q2 the FCA PSD have been grown in line with the SML data.

In this scenario we also assume that banks loosen their underwriting standards, permitting an increase in mortgages with LTI ratios up to 6, and performing affordability tests using a stressed rate of 5.5% instead of the 6.75% prevalent in 2016 (Table 1). This relaxation of policies is consistent with the broader scenario in which rising house prices, and the resulting stretch to borrower affordability, could reasonably be expected to cause some banks to adjust their affordability criteria in order to avoid a reduction in mortgage approvals or a loss of market share.

It is worth noting that neither scenario includes an additional deterioration in borrower quality above and beyond what is set in the underwriting standards above. For example, we do not assume any reduction in income for a fraction of new borrowers, described in Section 2.3, as we do not attempt to capture a pool of borrowers that is currently rejected but may not be if the underwriting standards were to loosen.

In the hypothetical upside scenario, macroprudential policies could have a substantial impact by excluding borrowers that, as a result of the rapid increase in house prices, are not able to obtain a mortgage with a loan value that is many times their income. This could reduce mortgage approvals, and by excluding borrowers that would need higher levels of debt in order to acquire property, the overall level of mortgage debt. The resulting LTI and DSR distributions would be skewed towards lower multiples relative to a counterfactual where the macroprudential policies are not in place.

4.1. Central scenario: Flow

Figure 6 shows the difference between the projected 2023 Q4 distributions of new mortgages (with and without policy) against the observed distribution of new mortgages in 2016 Q4.

As expected, there is very little change to the flow of new mortgages in the central scenario. The moderate increase in the house price to income ratio causes the LTI distribution to shift to the right. With macroprudential policy in place, this distribution is pushed to the left, and there is greater bunching just below the macroprudential flow limit of 4.5 LTI. This is the result of some borrowers being rejected a loan and others adjusting by finding cheaper properties and taking smaller loans, to avoid being denied a loan. A similar impact is seen on the DSR distribution, with the distribution shifting slightly to the right (as a direct result of higher LTI ratios). The size of this shift is also limited by the presence of macroprudential policies, but not substantially. The policy impact on approvals would be small in this hypothetical scenario, only about 3% of mortgages would be rejected relative to a no-policy outcome.

The LTV distribution shifts slightly right over the scenario as a result of household deposits growing more slowly than house prices. The effect is not substantial because increases in property prices, even for large differences in house price and deposit growth, do not increase LTVs appreciably. For example, if a mortgage with an initial LTV of 80% experiences a 20% increase in property price, and the deposit stays the same, its new LTV will be just 83%. Consequently, large shifts in the LTV distribution are unlikely to be driven by rapid house prices growth alone. Instead they require a relaxation in banks' underwriting standards, specifically an increased willingness to extend loans to borrowers with less money for a deposit.

Figure 6. Key distributions for the flow of new lending in the <u>central</u> scenario

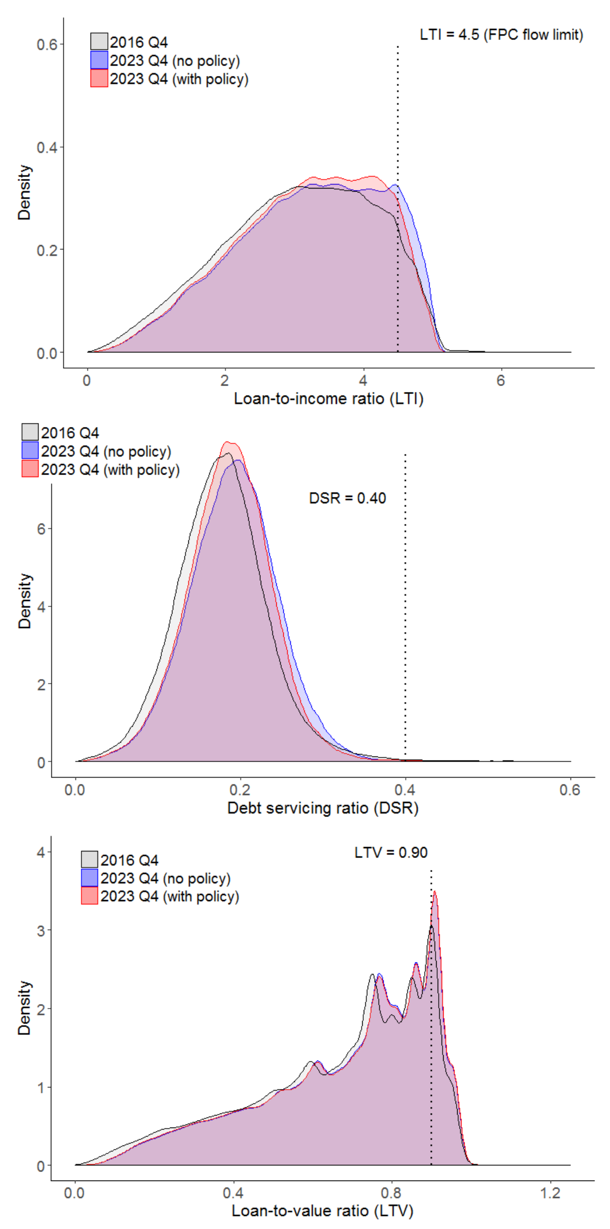

Figure 7. Key distributions for the flow of new lending in the <u>upside</u> scenario

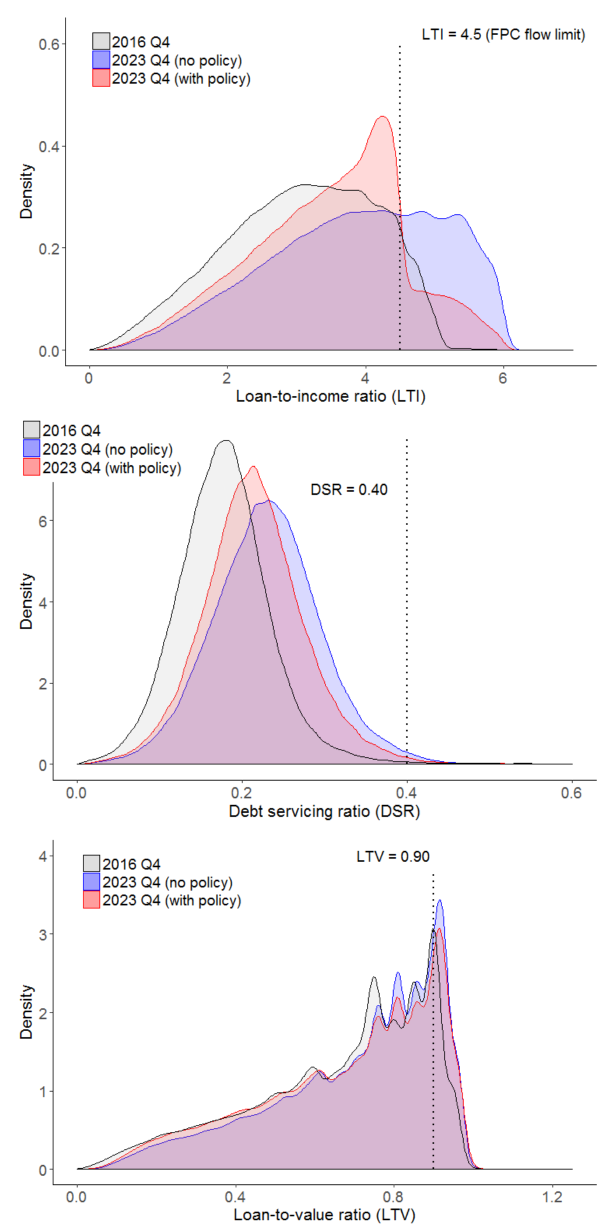

4.2. Upside scenario: Flow

The distributions for the upside scenario are shown in Figure 7. Relative to the central case, the LTI distribution shifts much more to the right, if macroprudential policies were not in place. This result is intuitive, given that house price growth significantly exceeds income growth in this scenario (cumulative growth of 48% for house prices compared to 24% for incomes).

In this case the LTI flow limit – permitting just 15% of mortgages to be extended at an LTI ratio at or above 4.5 – significantly impacts the flow of mortgages at the end of this hypothetical upside scenario. The share of lending above this threshold falls considerably, compared to a no-policy outcome. And there is substantial bunching just below the limit. As in the central case, this is the result of higher house prices making more borrowers require high LTI loans. In cases where a borrower's required LTI is above the threshold, and banks have already used up their 15% allowance, this borrower is either rejected a loan or forced to find cheaper housing. Because the policy is significantly more binding in this scenario, the cut-off just after the policy threshold is more dramatic. Overall, macroprudential policy reduces mortgage approvals by 13% in this upside scenario, compared to only 3% in the central scenario.

The rightward shift in the DSR distribution is almost entirely the result of changes in the LTI distribution, as the interest rate paths in the scenarios are identical. However, the change in the DSR distribution is moderated somewhat by borrowers extending mortgage terms in order to pass affordability criteria.

There is relatively little change in the LTV distribution, for the same reasons as in the central scenario – for a fixed (or slowly growing) deposit value, even large changes in property prices do not change LTVs substantially. In order to shift the LTV distribution more visibly, borrowers' deposits need to decrease substantially relative to property values.

The projected flows form a substantial part of what could happen to the UK mortgage stock under different macro scenarios. But in order to complete the picture, we need to add in the existing stock and current flow.

4.3 Central scenario: Stock

With little change in the flow of new mortgage lending, the stock of mortgages does not change much in the central scenario. So the impact of policy – which only binds marginally on the flow even at the end of the scenario – is diluted further when looking at the entire mortgage stock. The projected distributions for 2023 are presented in Figure 8.

The most notable change in the stock is the widening in the LTI distribution. Considering the development of the UK mortgage stock, both over the scenario and leading up to it, this is intuitive. The LTI distribution has shifted to the right relative to the early and mid-2000s. Thus, despite LTIs in the flow only increasing slightly over scenario, the repayment of older mortgages with lower origination LTIs and their replacement with new mortgages with higher origination LTIs produces this effect.

The LTV distribution shifts slightly to the left, which is also consistent with the recent history of the UK mortgage market. As mortgages in the current stock are gradually repaid, and as house prices increase, the indexed LTV ratio on these properties naturally declines, improving the left tail of the LTV distribution. But the right tail of the distribution remains broadly unchanged, as the improvements in the stock are offset by high origination LTVs in the projected flow.

4.4. Upside scenario: Stock

The main difference between the upside and the central scenarios is the substantially increased house price growth and approvals profiles. Aggregate income and unemployment are the same in the two scenarios, meaning the propensity of borrowers to suffer severe income shocks that push them into the tail of the LTI or DSR distributions remains constant.

Because of that, the variation in the projected stock between the two scenarios can be attributed to higher house prices or increased approvals. The impact of increased house prices on indexed LTVs will push older mortgages further down the LTV distribution than in the central scenario. And greater numbers of approvals – particularly remortgages and home movers – increase the relative share of the projected mortgage stock made up of mortgages in the projected flow. Both these effects are evident in Figure 9.

The general widening of the LTI distribution and shift to the left in the LTV distribution, evident in the central scenario, are also present here but to a greater extent. Both shifts are accelerated by the increase in house prices (which pushes LTVs further down, and forces new mortgagors to obtain higher LTI mortgages) and the number of approvals (which weights the final stock distribution closer to those of the projected flow).

Figure 8. Key distributions for the stock of mortgages in the <u>central</u> scenario

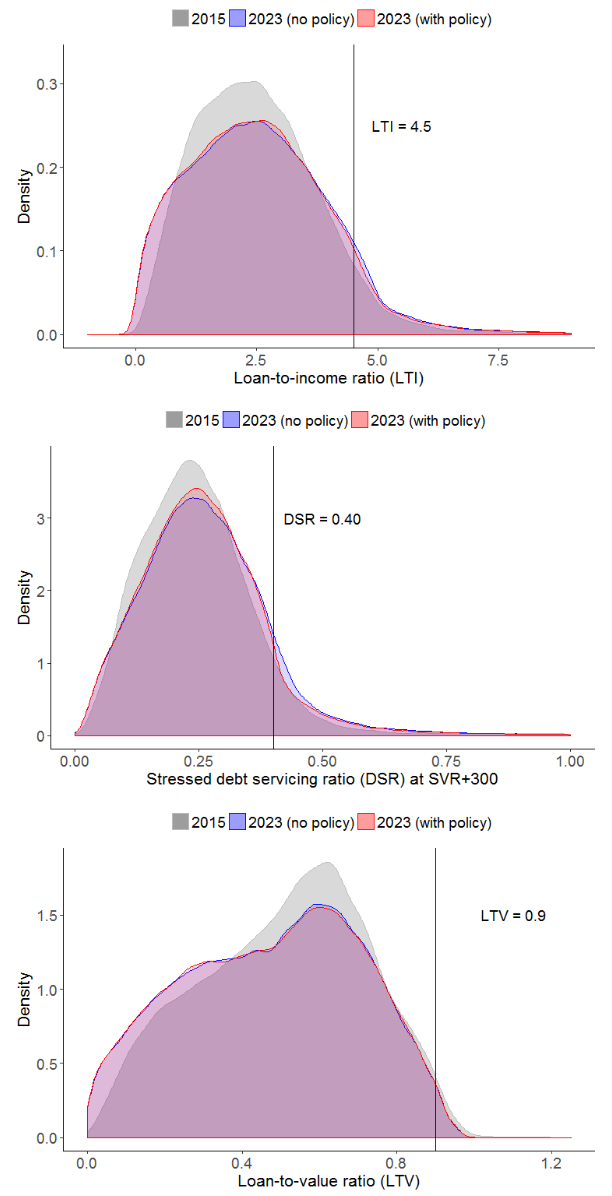

Figure 9. Key distributions for the stock of mortgages in the <u>upside</u> scenario

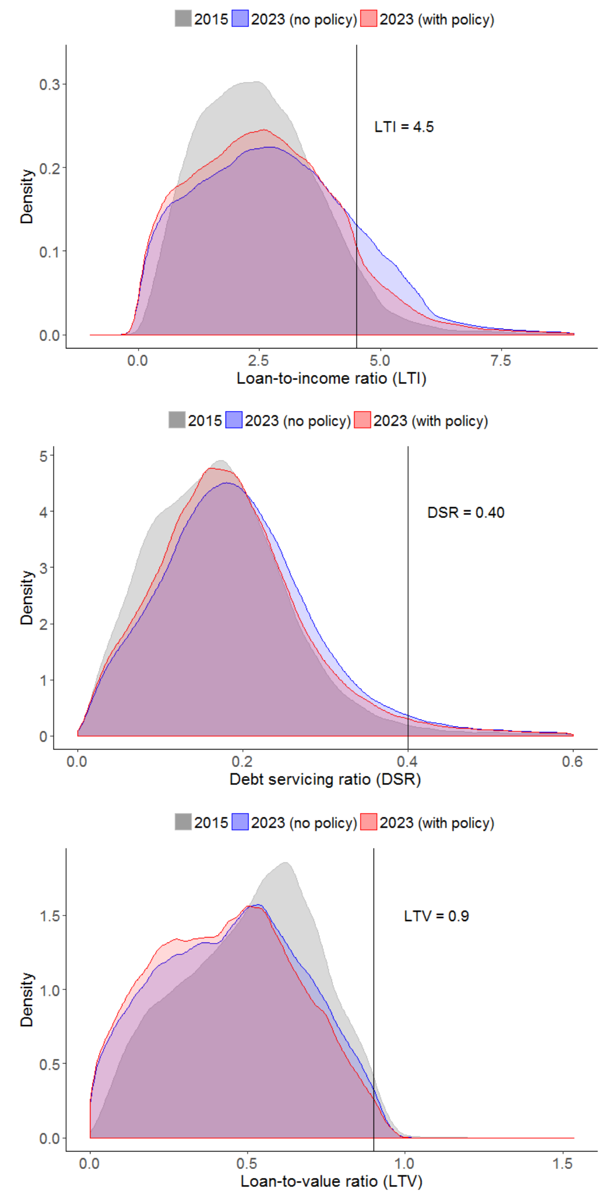

Macroprudential policy also has a bigger impact in the upside scenario than in the central case, by preventing the tail of highly indebted households from deteriorating as much as it would have in the absence of policy. In our hypothetical upside scenario the policy could reduce the share of mortgage stock with LTI at or above 4.5 by almost 6pp to 10.1%, compared to a less than 1pp reduction in the central case.

5. Conclusion

The approach outlined in this paper produces an estimate of how mortgage debt might evolve under different economic scenarios. This allows policy makers to assess the potential risks coming from household debt and to inform macroprudential policies.

We have shown that the vulnerabilities coming from household debt vary with the economic environment and that macroprudential housing policies can be used to mitigate them.

It is not the purpose of this paper to discuss optimal mortgage market policy, or the full set of costs and benefits associated with intervening in the mortgage market. However, the modelling framework outlined here, and associated outputs, provides policymakers with a set of forward-looking analytical outputs that can be used as part of the evidence base for assessing risks from household indebtedness and for deciding how to use macroprudential policies.

We will continue to improve the model and underlying assumptions. One of the challenges to forecasting the mortgage stock and its characteristics is accounting for remortgage activity. Our estimates could be improved by developing an approach that could identify potential remortgagors, based on the borrower characteristics and macroeconomic scenarios. Another challenge is forecasting potential income shocks, and more work can be done to improve our estimates. Finally, the model assumptions could be improved with more information about the behaviour of prospective borrowers and the adjustments they make, if they do not meet lenders' criteria.

References

Bank of England (2017), Financial Stability Report, June 2017, Issue No. 41.

Chakraborty, Chiranjit, Mariana Gimpelewicz and Arzu Uluc (2017), A tiger by the tail: estimating the UK mortgage market vulnerabilities from loan-level data, Bank of England Working Paper no. 703.

IMF (2018), The IMF's Annual Macroprudential Policy Survey – Objectives, Design, and Country Responses.

Jácome, Luis I. and Srobona Mitra (2015), LTV and DTI Limits – Going Granular, *IMF Working Paper*.

Jorda, Oscar, Moritz Schularick, and Alan M. Taylor (2016), The Great Mortgaging: Housing Finance, Crises, and Business Cycles, *Economic Policy*, Vol. 31 (85), pp. 107-152.

Lozej, Matija and Ansgar Rannenberg (2018), The Macroeconomic Effects of the LTV and LTI ratios in Ireland, *Applied Economics Letters*, Vo. 25 (21).

Mian, Atif, and Amir Sufi (2010), Household Leverage and the Recession of 2007-09, *IMF Economic Review*, Vol. 58 (1), pp. 74-117.

Reserve Bank of New Zealand (2017), Serviceability Restrictions as a Potential Macroprudential Tool in New Zealand, *Consultation Paper*.

Zabai, Anna (2017), Household Debt: Recent Developments and Challenges, *BIS Quarterly Review*.