

Heterogeneous Effects and Spillovers of Macroprudential Policy in an Agent-Based Model of the UK Housing Market

Adrian Carro^{*1,2}, Marc Hinterschweiger^{†3}, Arzu Uluc^{‡3}, and J. Doyne
Farmer^{§2,4,5}

¹*Banco de España*

²*Institute for New Economic Thinking (INET) at the Oxford Martin School, University of Oxford*

³*Bank of England*

⁴*Mathematical Institute, University of Oxford*

⁵*Santa Fe Institute, New Mexico*

July 26, 2022

Abstract

We develop an agent-based model of the UK housing market to study the impact of macroprudential policy experiments on key housing market indicators. The heterogeneous nature of this model enables us to assess the effects of such experiments on the housing, rental and mortgage markets not only in the aggregate, but also at the level of individual households and sub-segments, such as first-time buyers, homeowners, buy-to-let investors, and renters. This approach can therefore offer a broad picture of the disaggregated effects of financial stability policies. The model is calibrated using a large selection of micro-data,

^{*}adrian.carro@bde.es

[†]marc.hinterschweiger@bankofengland.co.uk

[‡]arzu.uluc@bankofengland.co.uk

[§]doyne.farmer@inet.ox.ac.uk

including data from a leading UK real estate online search engine as well as loan-level regulatory data. With a series of comparative statics exercises, we investigate the impact of (i) a hard loan-to-value limit, and (ii) a soft loan-to-income limit, allowing for a limited share of unconstrained new mortgages. We find that, first, these experiments tend to mitigate the house price cycle by reducing credit availability and therefore leverage. Second, an experiment targeting a specific risk measure may also affect other risk metrics, thus necessitating a careful calibration of the policy to achieve a given reduction in risk. Third, experiments targeting the owner-occupier housing market can spill over to the rental sector, as a compositional shift in home ownership from owner-occupiers to buy-to-let investors affects both the supply of and demand for rental properties.

JEL Classification: D1, D31, E58, G51, R21, R31

Keywords: Agent-based model, housing market, macroprudential policy, borrower-based measures, buy-to-let sector.

1 Introduction

Housing is an integral part of the economy. It is often the single largest asset of households and can substantially affect their saving and consumption decisions. In many countries, mortgage finance is one of the most important business lines of the financial sector. From a macroeconomic perspective, financial crises that are preceded by credit booms in the housing market tend to be associated with deeper recessions and slower recoveries ([Jordà et al., 2016](#)). Yet a housing sector is often missing in canonical macroeconomic models, or represented in a simplistic manner. Depending on the research question of interest, this may obscure important dynamics in the economy and may lead to misleading policy prescriptions. Furthermore, it is difficult to assess the effects of policy interventions on different segments of the housing market, such as renters, first time buyers (FTBs), home movers (HMs) or buy-to-let (BTL) investors, unless these segments and their corresponding dynamics are properly considered within the model.

In this paper, we develop an agent-based model (ABM) of the housing market, which allows us to overcome some of these limitations and paint a more detailed picture of the different segments of the housing market, their dynamic properties, and their interactions with each other. We use the model to assess the effects of macroprudential housing policy experiments on housing and mortgage variables, with a particular focus on their time-series and cross-sectional behaviour. The model has several types of households (renters, owner-occupiers such as FTBs and HMs, and BTL investors) who are subject to a life cycle, which influences their propensity and ability to transact in the housing market. For example, a new household would typically start out by renting a dwelling, may buy its first property once it has saved a sufficiently large deposit, and it might later invest in additional properties to rent out. This allows for realistic and detailed housing market dynamics, including endogenous house price cycles and households whose type changes endogenously.

A preliminary version of this model has been previously presented in [Baptista et al. \(2016\)](#). With respect to that earlier version, this paper extends the model, analysis and discussion in different directions. First, it provides a more accurate and detailed description of the model. Second, the design of the model has been greatly improved in a number of key aspects: (i) many of the existing mechanisms have been enhanced, such as the demographics block,

which is now able to preserve the shape of the age distribution; (ii) a few new mechanisms have been implemented, such as age requirements for mortgage borrowing; and (iii) several model components less central to our research focus have been simplified or removed, such as interest payments on bank deposits. Third, while a large fraction of parameters were manually set—as opposed to empirically estimated—or even undisclosed in the earlier version, we have now implemented a systematic empirical estimation and calibration of the full set of model parameters. Importantly, this combination of model modifications and better estimation and calibration of parameters has led to a substantial improvement of the corresponding output, bringing it much closer to empirically observed values in multiple dimensions. Fourth and finally, we have simulated a new set of experiments and we include here a much more detailed analysis of results.

We calibrate the model to the UK housing market and use it to investigate the impact of housing tools available within the UK macroprudential framework. In the aftermath of the Global Financial Crisis, the Financial Policy Committee (FPC) within the Bank of England was formed as a macroprudential regulation authority in the UK. The FPC was granted powers to address risks to financial stability arising from the housing market. These powers allow the FPC to restrict the proportion of new mortgages by introducing limits on loan-to-value (LTV) and debt-to-income (DTI) ratios in the owner-occupier mortgage market, and LTV and interest coverage ratios with respect to buy-to-let lending. The LTV tool is related to credit risk and aims to protect lenders' capital in the event of default. On the other hand, the DTI tool¹ is related to direct risks to the financial system due to defaults resulting from unexpected decreases in income or increases in interest rates or indirect risks to the wider economy due to a reduction in consumption in order to continue servicing debts. In October 2014, the FPC introduced a loan-to-income (LTI) flow limit and an affordability test to insure against loosening in underwriting standards.² The aim of this paper is to investigate the potential impacts of housing tools available within UK's macroprudential framework but not to evaluate the rules currently in place.

We use the model to perform a series of comparative statics exercises to investigate the

¹ A DTI tool applies to both a borrower's mortgage and non-mortgage debt. A loan-to-income (LTI) tool can be regarded as a version of a DTI tool that applies only to a borrower's mortgage debt.

² For further information, see www.bankofengland.co.uk/record/2014/financial-policy-committee-june-2014.

impact of two of these housing policy experiments: (i) a hard LTV limit, and (ii) a soft LTI limit, allowing a certain share of new mortgages above this limit. We assess how these policy experiments affect key variables of interest, such as the number of housing transactions, LTV and LTI ratios, house prices and mortgage approvals, among others. Our agent-based approach allows us to examine the joint distribution of multiple housing market variables of interest at the household level. We can therefore assess the aggregate dynamics in the UK housing market, and at the same time identify the key drivers and most affected sub-groups of agents at a micro level.

We find that, first, the housing policy experiments tend to mitigate the house price cycle by reducing credit availability and therefore leverage. Second, a policy experiment targeting a specific risk measure (e.g. lending at high LTI ratios) may also affect other risk metrics (such as LTV ratios). Therefore, when calibrating a policy to achieve a given desired reduction in risk, the effects of the policy on all risk metrics should be taken into account. Third, policy experiments targeting the owner-occupier housing market benefit BTL investors, who can increase the number of houses they own. Among owner-occupiers, FTBs are more affected than HMs, facing a more limited access to high-risk mortgages. These changes also spill over to the rental sector, as a compositional shift in home ownership from owner-occupiers to BTL investors affects both the supply of and demand for rental properties.

The paper is structured as follows. The next section briefly reviews the literature in the context of our model. Section 3 describes the model. Section 4 sets out the validation of the model. Section 5 reports the results of our policy experiments. Section 6 concludes.

2 Literature Review

In the aftermath of the Global Financial Crisis, the implications of heterogeneity for macroeconomics have become a more central topic for research. For example, [Kaplan et al. \(2018\)](#) introduce a continuum of heterogeneous households in a canonical DSGE model to assess the transmission mechanism from monetary policy to household consumption. But several studies have highlighted the inherent limitations of the mainstream approach for adequately capturing heterogeneity, and instead propose agent-based models as a viable alternative (see e.g. [Farmer](#)

& Foley, 2009; Fagiolo & Roventini, 2017; Dosi & Roventini, 2019; Haldane & Turrell, 2019).

Over the last decade or so, several macroeconomic agent-based model (ABM) “families” have been developed. They often assess similar questions as mainstream models, but with a stronger focus on the implications of heterogeneity for macroeconomic dynamics as well as on the impact on different types of agents. We broadly follow the grouping described in the overview article by Dawid & Delli Gatti (2018). Dosi et al. (2010) have developed the “Keynes meeting Schumpeter” family, which amalgamates features of both demand-side Keynesian economics and economic growth following Schumpeter. A second family is based on Complex Adaptive Trivial Systems (CATS), as set out by Delli Gatti et al. (2011). Several notable extensions include Assenza et al. (2015), who introduce capital goods, and Caiani et al. (2016), who focus on stock-flow consistency. A third family of models, JAMEL (Java Agent-based MacroEconomic Laboratory), focuses also on stock-flow consistency as well as on the learning and adaptive behaviour of agents (Seppecher et al., 2018, 2019). Fourth, Ashraf et al. (2017)’s model explores the role that banks play in a network of trading firms, and the effects of microprudential regulation and financial stability policy. The EURACE family comprises a large range of different strands, most importantly the Genoa branch (Cincotti et al., 2012a) and the Bielefeld branch (Dawid et al., 2012, 2014, 2019). Finally, the LAGOM model family was started by Haas & Jaeger (2005) to investigate the aggregate outcomes when various agents manage climate risks, and was subsequently developed into a fully-fledged macro ABM by Mandel et al. (2010) and Wolf et al. (2013).

There also exists a large number of ABMs tailored to specific issues. We briefly summarise below recent studies that are most closely linked to our research in this paper.

Many ABMs have been developed to generate non-linearities and complex dynamics in the housing market. The seminal work of Geanakoplos et al. (2012) and Axtell et al. (2014) shows that the US housing market bubble leading to the Global Financial Crisis was due mainly to changes in loan-to-value (LTV) ratios rather than interest rates, using a comprehensive dataset on the housing market in the Washington, D.C. metropolitan area. Similarly, Ge (2013) and Ge (2017) demonstrate that weak underwriting standards may result in a housing bubble. Glavatskiy et al. (2021) use a large-scale ABM to model house price dynamics in the

Greater Sydney region, explained by households' trend-following behaviour and ability to borrow. [Evans et al. \(2021\)](#) use a spatial ABM, also of the Greater Sydney region, to analyse the effects of social and economic factors on agents' decision-making behaviour. A related strand of the literature, represented by [Parker & Filatova \(2008\)](#), [Huang et al. \(2013\)](#) and [Filatova \(2015\)](#), combines spatial housing ABMs with land use models.

There are also several ABMs developed for the analysis of macroprudential housing policies. Based on [Baptista et al. \(2016\)](#), the studies of [Mérő & Vágó \(2018\)](#), [Cokayne \(2019\)](#), and [Catapano et al. \(2021\)](#) develop ABMs of the Hungarian, Danish, and Italian housing markets, respectively, and perform various macroprudential policy experiments. [Laliotis et al. \(2020\)](#) create an ABM for countries of the euro area to assess the effects of LTV limits. [Ozel et al. \(2019\)](#) incorporate a housing market in the EURACE model mentioned above and apply a household net worth target and debt-service-to-income limit to determine the impact on macroeconomic performance. [Yun & Moon \(2020\)](#) apply LTV and debt-to-income policy experiments to an ABM of the Korean housing market.

Another focus of housing ABMs is inequality. [Pangallo et al. \(2019\)](#) develop a spatial ABM to study the relationship between income inequality and house prices. [Guerrero \(2020\)](#) assesses the role of the decentralized nature of housing market transactions for the distribution of housing wealth in a one-to-one scale model of all households in the UK. Building on [Baptista et al. \(2016\)](#), [Tarne et al. \(2021\)](#) investigate the link between household debt, wealth inequality and consumption volatility across the housing market cycle, and which macroprudential tools are most effective in improving these macroeconomic outcomes.

Building upon [Geanakoplos et al. \(2012\)](#) and [Axtell et al. \(2014\)](#), this paper contributes to the literature by expanding existing housing ABMs along different dimensions. Most importantly, a combination of three key features set our model apart from previous contributions, including these two seminal works. First, we introduce a realistic life-cycle dynamics, including death, inheritance and plausible income trajectories, and allowing us to better match the distribution of several variables of macroprudential relevance. This life-cycle dynamics is essential for our policy experiments: as these policies are constraints on the flow of new lending, the life cycle ensures a realistic continuous demand for mortgages from borrowers as many of them

turn from renters into first-time buyers, home movers, and eventually bequeath their properties to the next generation. To the best of our knowledge, our agent-based model is the first to allow for such a flow limit, comprising both hard and soft caps. Second, we consider a more realistic modelling of market interactions: a double-auction market mechanism. This enables us to generate more realistic loan-to-income (LTI) distributions, which is particularly important for the soft LTI policy experiment we perform. Third, our model also includes a dynamic buy-to-let sector and a rental market, which allows us to trace spillovers from housing market policies targeting owner-occupiers on the rental market. Again, to the best of our knowledge, our ABM is the first to allow for a detailed modelling of such interactions between these sectors, and for an assessment of how macroprudential policy experiments affect these interactions within a single framework. Furthermore, the model has the potential to augment the toolkit available for policy-making at central banks, allowing them to exploit the heterogeneity and non-linear effects that are rarely captured in other models. This is especially relevant for macroprudential policy design as such policies often target sub-sectors of the economy, such as borrowers with certain characteristics. We have demonstrated the usefulness of this approach by jointly analyzing the effects of such policies on various risk metrics, such as loans with high LTV and LTI ratios. A final, technical, contribution is provided by the detailed empirical estimation of behavioural rules and parameters, leaving only a few of these parameters to be calibrated with the method of simulated moments.

3 Model description

There are three types of agents in the model: *(i)* households, *(ii)* a bank, and *(iii)* a Central Bank. Households are the main agents; they interact with each other by buying and selling houses in the sales market, as well as by renting and letting their investment properties out in the rental market. They also interact with the bank by requesting mortgages to buy houses, which the bank provides according to its own internal policy requirements. On top of this bank internal policy, the Central Bank further regulates the conditions under which mortgages can be underwritten.

At the beginning of each simulation the target number of households are created,³ drawing

³ While we have tested the model with up to two million households, due to the high computational requirements

all relevant household characteristics from suitable distributions (independently unless stated otherwise). Notably, households are assigned: (i) an initial age of the Household Reference Person,⁴ which will evolve in time and drive the evolution of other household attributes; (ii) a permanent income percentile, which will remain fixed throughout the simulation; (iii) a permanent propensity to save; and (iv) an initial financial wealth, dependent on the household's initial income and its propensity to save. In addition, a buy-to-let (BTL) flag is randomly added to some households, with a probability dependent on income percentile, signalling their capacity or willingness to invest in BTL properties. Households receiving the BTL flag are also assigned an investor type —capital-gains-driven, rental-income-driven or mixed—, associated with a specific intensity of their interest in capital gains as opposed to rental yield. Finally, all newly created households are initially set to live in social housing, which we define as a temporary accommodation, with no cost, while trying to find a house to rent or buy. In a broad sense, this concept of social housing can represent situations such as homelessness, living with parents or living out of housing and other social benefits. Even though no housing payment is deducted from the households' income while in social housing, they (i) never choose to be in that position, (ii) always try to secure a different accommodation option, and (iii) are only put there if they fail to secure any other form of housing or as a temporary state between selling a house and buying or renting a new one.

In a similar manner, the number of houses corresponding to the target ratio of houses to household are also created at the beginning of the simulation and distributed at random among existing households,⁵ with this housing stock remaining constant thereafter. Since all households are initially in social housing, those receiving a house will simply move into it. Each house is characterised by a single, discrete parameter, its quality, which is a proxy for all possible features differentiating houses from each other and making some more desirable than others, such as their location, size, condition or dwelling type. This discrete house quality is assigned

of the calibration process, all results reported here are for 10,000 households.

⁴ The concept of Household Reference Person is used in a number of government surveys for deriving statistics and describing the household in terms of the characteristics of the household member selected for the purpose of a survey. For a precise definition see Appendix A.2.

⁵ A random distribution of houses to households at the beginning of the simulation may seem crude. Our intention is to avoid imposing any structure or correlation between income and financial wealth, and home ownership and quality, and instead allow the model to generate that structure endogenously —which it successfully does after a few hundred time steps.

at random at initialisation and remains fixed during the simulation, thus dividing the market into equally sized segments. Importantly, this reduction of all house features into a single dimension allows for all households to have the exact same order of preferences regarding a given set of houses. For further details about the initial set-up of the model see Appendix A.1.

The following subsections provide a general description of the model, leaving some of the details for Appendix A. Throughout these subsections, we use Latin letters for identifying specific model variables and Greek letters for parameters, with the exception of ε , used to denote a random draw from a normal distribution with mean ε_μ and standard deviation ε_σ , η , used to denote a random draw from a distribution estimated from data, and σ , used to denote the standard logistic function $\sigma(x) = 1/(1 + e^{-x})$. The specific values of the parameters, as well as a description of the estimation or calibration procedure and the data sources used, can be found in Tables 8 to 14 in Appendix B by referring to the corresponding subsection and/or equation number. Further details on estimation and calibration procedures can be found in the Online Appendix.

3.1 Overview of the simulation step

Once the model has been initialised as described above (see also Appendix A.1), it proceeds iteratively in simulation steps of one month, where the following updates of model components occur:

1. **Demographics:** New households are born, some households die and the rest of them age. Since the ageing process alone would modify the age distribution, we choose both the birth and the death rate —both dependent on age— so as to counteract this effect and thus keep the age distribution constant, as well as to keep the total number of households fluctuating around its target. As those created at model initialisation (see above), new households are born with all their characteristics —initial age, income percentile, propensity to save, initial financial wealth and, potentially, BTL flag and type— drawn from suitable distributions, and are set to live in social housing. When a household dies, another randomly chosen household inherits all financial and housing wealth. See Appendix A.2 for further details.

2. Households:

(a) Receive income, consume and pay housing expenses: Households receive a monthly income conditional on age, evaluated at the household's income percentile. Thus, the household's income changes along its lifetime, driven by age, even if its income percentile remains fixed. Additionally, households with BTL investments collect rent, if applicable. They pay income tax and make National Insurance contributions, and pay for their essential non-housing consumption. Then they pay their housing expenses: Renters pay rent, and owners with a mortgage make their mortgage payments. After meeting all their monthly financial commitments and essential expenses, households decide on their additional desired consumption, as a function of their remaining financial wealth. For further details, see Subsection 3.2.

(b) Make their housing decisions, according to their current status:

- If in social housing, they decide whether to rent or buy a new house;
- If renting, they continue to rent, that is, tenants never decide to leave a rental property before the contract expires;
- If owner-occupying a house and not having the BTL flag, they decide whether to sell their house;
- If owner-occupying a house and having the BTL flag, they decide whether to buy a new investment property and, for each vacant investment property owned, they decide whether to sell it.

Apart from this, house owners with unsuccessful offers from previous months, whether on the sales or the rental market, decide whether to update the price of these offers.

For further details, see Subsection 3.3.

(c) Place their bids and offers on the relevant market, whether sales or rental.

3. Bank: Provides mortgages as requested by households, complying both with its internal lending policy as well as with that imposed by the Central Bank. After all transactions have been processed and all mortgages signed, the bank updates its mortgage interest rate

for next month, based on the rate for this month and the resulting excess demand for credit over the bank's target. For further details, see Subsection 3.4.

- 4. Markets:** The sales market is cleared first, then the rental market, so that BTL investors can directly offer their newly acquired investment properties for rent. Both markets clear in the same way, matching bids and offers in a number of rounds until no match is possible. In each of these rounds, first, bidders are matched to the best offer they can afford with their bid price. Then, offering households select one of these matched bids at random, after potentially increasing their offer price in case they have received multiple bids. Unsuccessful bids are dropped at the end of the clearing process, while unsuccessful offers are kept for the following month. For further details, see Subsection 3.5.

3.2 Income, non-housing consumption and financial wealth

A household's gross monthly income is modelled as a function of its age and income percentile. In particular, we obtain a distribution of gross income conditional on age from data, which we then use to compute a household's gross income as the inverse cumulative distribution function corresponding to the household's age and evaluated at the household's income percentile. In this way, as households age, their income percentiles remain fixed but their actual incomes change to reflect the income distribution over households of that age. The fixed income percentile assumption effectively implies that there are no idiosyncratic shocks, such as unemployment. Note that the income generated in this way does not include rental income, which is endogenous to the model and thus explicitly accounted for. Therefore, for households with BTL investments, their current gross rental income has to be added in order to arrive at their gross total income.

This gross total income is subject to tax payments in the form of an income tax⁶ and National Insurance contributions, with no capital gains tax applied at any point in the model. Furthermore, households are assumed to have an unavoidable minimum level of monthly consumption, which we call essential non-housing consumption and estimate with reference to consumption levels among households with incomes around the minimum government income

⁶ When computing the income tax liability, we take into account the tax relief available to BTL investors with rented out properties.

support for a married couple.⁷

Finally, after having met all their monthly financial commitments and essential expenses, including housing expenses, whether rental or mortgage payments⁸, households decide on their additional desired consumption, i.e. on their non-essential non-housing consumption. It is important to recall at this point that, with the model lacking any macroeconomic dynamics, the only role of non-housing consumption is to determine the households' financial wealth, which is, in turn, of utmost importance for determining the down-payments they can make when buying houses and thus mortgages available to them. In this context, we model household consumption in such a way that leads to an accurate distribution of financial wealth among the population. To this end, we first define, for each household, a target or desired level of financial wealth using a distribution of financial wealth conditional on gross income and the household's propensity to save and then set its desired consumption to be a fraction of the difference between the current financial wealth of the household and its desired level (see Appendix A.3 for further details). In this way, the financial wealth of each household relaxes exponentially towards its desired level, thus ensuring that the distribution of financial wealth in the model remains close to the one observed in the data. Finally, households save any remaining part of their income, or draw down their financial wealth if consumption is larger than their disposable income after all previously described expenses.

3.3 Housing decisions

Every month, households make their housing decisions depending on their current housing status: households without a home (in social housing) decide whether to rent or buy a new house, homeowners without the BTL flag decide whether to sell their home, and homeowners

⁷ While these expenses are unavoidable, the specific parameter values used lead to an essential non-housing consumption smaller than the minimum possible net income, and are thereby unable to lead to bankruptcies.

⁸ It is possible at this point that a household cannot afford these payments. Since the model does not seek to capture the nuances of bankruptcy dynamics (from delinquency to foreclosure), we simply consider these households as directly bankrupt. In this case, we ensure scheduled payments are always made by artificially injecting as much cash as necessary, without further action against the bankrupt household (see Appendix A.4 for further details). Note that in the UK, the number of arrears and defaults has been relatively low during the time period under consideration, and does not appear to drive housing market dynamics. Also, ad hoc government policies may support homeowners during times of crisis, for example mortgage payment holidays during the Covid-19 pandemic.

with the BTL flag decide whether to increase or decrease their investment portfolio. Renters simply continue to rent until the end of their current rental contracts, when they make the same decisions as if they were in social housing.⁹ Thus, no specific subsection is devoted to them here.

3.3.1 Decisions if in social housing

All households in need of a new home are represented as being in social housing, including newly formed households, previous renters whose contracts have expired, and previous homeowners having sold their homes. These households have to decide *(i)* whether to rent or buy a new house, i.e. in which market to bid, *(ii)* how much to bid in the chosen market, and *(iii)* how much to pay as a down-payment in case they decide to buy. In order to make the first decision, they first decide on a desired purchase price, i.e., on how much they would like to spend on buying a house, given the maximum mortgage available to them. Then, they find out the maximum house quality they can afford with that desired purchase price. Afterwards, the decision is made by comparing the annual costs of buying for that desired purchase price with those of renting a house of the same quality. If they decide to buy a house, they simply go ahead with the computed desired purchase price, while they compute a specific desired rental price if the decision is to rent. Finally, if deciding to buy, households must also choose a desired down-payment, which might later on be overruled if it is below the minimum down-payment required by the bank.

In the spirit of [Axtell et al. \(2014\)](#), we model the desired purchase price p_d as an exponentially noisy and nonlinear function of the household's gross annual income, capped by the price corresponding to the maximum mortgage available to the household p_{\max} ,

$$p_d = \min(\alpha y^\beta e^\varepsilon, p_{\max}) , \quad (1)$$

where y is the household's gross annual income and ε is a normal noise. However, as opposed to the original specification by [Axtell et al. \(2014\)](#), we do not consider a denominator depen-

⁹ Importantly, most households will manage to secure a new accommodation starting right after the end of their previous tenancy. Only those failing to do so will actually spend time in social housing.

dent on the current price trend, as we did not find this component to significantly improve the estimation with our data. Furthermore, for the sake of the decision between renting and buying, we consider a version of this desired purchase price capped by prime quality prices,

$$p'_d = \min(p_d, \overline{p_{Q_{\max}}}) , \quad (2)$$

where $\overline{p_{Q_{\max}}}$ is the exponential moving average sale price of houses of the maximum quality (see Appendix A.7 for details on the market price information available to households).

The purpose of this latter cap is to prevent comparing artificially large desired purchase prices, potentially detached from current market dynamics, with rental prices which are by definition (see below) based on current market dynamics and thus unlikely to be unreasonably large.

The annual cost of buying a house for the household's desired purchase price capped by prime quality prices, p'_d , can be defined as the annual mortgage payments, $12m$, where m is the monthly payment of the mortgage required by the household for such a purchase, minus the annual house price appreciation, $p'_d g$, where g is the expected annual house price growth. Based on survey data (see Appendix B Table 10 for details), we model households' house price growth expectations as a linear function of past house price growth,

$$g = \gamma \text{HPA} + \zeta , \quad (3)$$

where HPA is the (geometric) mean annual house price appreciation between the most recent quarter and the quarter two years before. In order to compare this annual purchase cost with an equivalent rental option, the household must first find out the house quality Q they can afford to buy under current market conditions. They do this by comparing this desired price with the exponential moving average sale prices of the different house qualities. Then, the annual cost of renting is computed as the annual exponential moving average rent price for a house of this quality Q under current market conditions, r_Q , multiplied by a factor $(1 + \lambda)$, which represents a psychological cost of renting¹⁰. The total annual cost of buying is $12m - p'_d g$, while that of

¹⁰ We assume that, given the same overall cost, households inherently prefer buying a property over renting it, if

renting is $r_Q(1 + \lambda)$. The discrete choice between these two options, buying versus renting, can be modelled as a logistic function of the difference between their respective costs, with a sensitivity parameter θ setting how noisy or deterministic the decision is. In this way, the probability of deciding to buy, as oppose to rent, can be written¹¹ as

$$P_{\text{buy}} = \sigma\left(\theta[r_Q(1 + \lambda) - (12m - p'_d g)]\right), \quad (4)$$

where the logistic function is given by $\sigma(x) = 1/(1 + e^{-x})$. As an exception to this decision rule, households with the BTL flag will always decide to buy, in order to avoid situations where BTL investors themselves are renting or bidding in the rental market.

In case the household decides to buy a house, it simply sends a bid to the sales market for its desired purchase price p_d , defined in Eq. (1). Equivalently, if the household decides to rent a house, then it sends a bid to the rental market for its desired rental price r_d . Similarly to its purchase price equivalent, the desired rental price r_d is also defined as a nonlinear function of the household's gross annual income y ,

$$r_d = \min(\mu y^\nu, r_{\max}), \quad (5)$$

but in this case the function is not noisy, reflecting the far smaller fluctuations in desired rental prices as compared to desired purchase prices, and it is capped by the available net income of the household after essential consumption, r_{\max} .

If a household chooses to buy, then it decides whether to buy outright with cash or by obtaining a mortgage. If the financial wealth of the household is greater than the price of the house, it will pay for the house in cash,¹² otherwise, it will apply for a mortgage and choose a desired down-payment. While a minimum down-payment is determined by the mortgage

they had either option. This 'psychological cost of renting' factor determines how much lower the cost of a property, if rented, would need to be in order for households to be indifferent between renting or buying it.

¹¹ We followed the theory of the user cost of owner-occupied housing and the rental equivalence, which was first developed by Gillingham (1980), in designing the decision of buying vs renting. Gillingham (1980) defines the user cost of housing as "the opportunity cost of holding the house less the increase in the house's value". In equilibrium, this should be equal to the cost of renting for the same period of time.

¹² Note that this rule of behaviour, while extremely simple, is enough for the model to generate a roughly correct proportion of cash buyers.

conditions set by the bank, the household may choose to make a larger down-payment. In particular, since first-time buyers (FTBs) tend to use most of their financial wealth when buying a house, we model their desired down-payments as their full current financial wealth. For all home movers (HMs), a desired down-payment is computed using a distribution estimated from data by analysing the inverse cumulative distribution at the household's income percentile and multiplying this result by the current house price index, such that down-payments adapt to current market conditions.

3.3.2 Decisions as an owner-occupier

Households sell their owner-occupied property on average every 17 years, based on UK data, due to exogenous reasons not addressed in the model, such as starting a family or divorce. In order to prevent landlords from selling their homes and potentially becoming unable to buy a new one or even deciding to rent, we only allow homeowners without the BTL flag to sell their homes. If the household decides to sell its home, it will offer it on the market at a price p_s given by

$$\ln p_s = \ln(\overline{p_Q}) + \eta, \quad (6)$$

where $\overline{p_Q}$ is the exponential moving average sale price of houses of this quality and η is a random mark-up drawn from an estimated distribution (see Appendix B Table 11 for estimation details). This is a similar equation to the one used by [Axtell et al. \(2014\)](#), though simplified to focus only on the influence of average prices of similar houses, which we find to be the main driver in our estimation.

For each month a house remains unsold on the market, the seller has a certain probability of reducing its price. When a reduction takes place, the new price will be $p_s^t = p_s^{t-1} (1 - \exp(\varepsilon))$, where t indicates the time step and ε is a normal noise. Both the probability of a price reduction and the specific mechanism or functional form of the updated price are estimated from data (see Appendix B Table 11 for estimation details). If the seller has a mortgage on this property and the price drops below the corresponding outstanding mortgage principal owed to the bank, then the offer is withdrawn from the market. After successfully selling its home, a household is temporarily set to be in social housing, where it will have to decide again whether to buy a new

home or switch to renting.¹³

3.3.3 Decisions as a BTL investor

Each month BTL investors decide whether to bid for a new investment property, and for each owned investment property which is vacant or at the end of its tenancy agreement, they decide whether to keep it or sell it. Both these decisions are based on the expected yield to be obtained from such investments, which is a function of the expected rental yield and the expected capital gains. This set-up is designed to capture the various motives that investors may have, such as obtaining a steady stream of rental income or seeking capital gains to be realised only after the sale of the property. In particular, we consider three types of investors with different weights on capital gains as opposed to rental yield: (i) capital-gain-driven investors, with a large weight on capital gains; (ii) rental-income-driven investors, with a small weight on capital gains; and (iii) mixed investors, with equal weights on capital gains and rental yield. Moreover, we use survey data to estimate the respective shares of these types of investors (see Appendix B Table 12 for estimation details).

When deciding whether to bid for a new investment property, BTL investors compute the expected yield V_{buy} of a hypothetical house, making use of the highest amount of leverage available to them. Leverage is computed as the price corresponding to the maximum mortgage available to the investors, p_{max} , divided by the minimum down-payment the investors could make for such a mortgage, d_{min} . Moreover, they assume the expected rental yield of this hypothetical investment property to be the current exponential moving average rental yield over all house qualities \bar{s} , which takes into account the expected average occupancy of an investment property. Finally, they compute the mortgage cost as the annual mortgage payment corresponding to the maximum mortgage available to the household, $12 m_{\text{max}}$, divided by the initial equity on the investment, which is again the minimum down-payment d_{min} . In this way, the expected

¹³ While in reality households would try to secure an alternative accommodation before selling their home —e.g., via short-term rentals or by coordinating the sale and purchase to happen simultaneously—, a detailed depiction of these mechanisms is beyond the scope of our model, which focuses only on those aspects of the housing market most likely to matter for macroprudential policy. Given the larger number of rental as compared to sales transactions, implementing short-term rentals for home movers in our model would at most have a negligible impact on prices and tenancies. The coordination of simultaneous sales and purchases, on the other hand, would play the role of a market friction delaying transactions, but with no effect on the balance between supply and demand, and thus no impact on prices.

yield can be written as a weighted sum of both rental and capital yield rates times the leverage minus the mortgage cost,

$$V_{\text{buy}} = \frac{p_{\text{max}}}{d_{\text{min}}} (\delta_i g + (1 - \delta_i) \bar{s}) - \frac{12 m_{\text{max}}}{d_{\text{min}}}, \quad (7)$$

where g is the expected annual house price growth, as set in Eq. (3), and δ_i is the weight given by the specific investor type i to capital gains as opposed to rental yield. By modelling the annual probability to buy a new investment property as a logistic function of its expected yield, we can write the corresponding monthly probability as

$$P_{\text{buy}}^{\text{BTL}} = 1 - (1 - \sigma(\xi V_{\text{buy}}))^{\frac{1}{12}}, \quad (8)$$

where ξ is a sensitivity parameter, setting how noisy or deterministic the decision is.¹⁴ If the investor decides to buy a new property, she will place a bid in the sales market for an amount corresponding to her maximum available mortgage, p_{max} . Finally, if the bid is successful, the investor will pay for the house outright if she has enough financial wealth to do so, requesting a mortgage from the bank otherwise. In case of requiring a mortgage, the investor will also choose a desired down-payment, which we model as a draw from a normal distribution of down-payment fractions with parameters estimated from data and thus independent of income considerations. This desired down-payment is further multiplied by the current house price index, such that investor down-payments adapt to current market conditions.

Investors always try to rent out their vacant investment properties, that is, as soon as a rental contract ends, or whenever a new investment property is bought or inherited, the owner offers it for rent on the rental market. Similarly to the ask price for offers in the sales market [see Eq. (6)], investors offer their investment properties in the rental market for a monthly rent r_s given by

$$\ln r_s = \ln(\overline{r_Q}) + \eta, \quad (9)$$

where $\overline{r_Q}$ is the exponential moving average rent price of houses of this quality and η is a

¹⁴ Note that the probability of buying this month is equal to one minus the probability of not buying this month, which is equal to the probability of not buying this year raised to the power of 1/12.

random mark-up drawn from an estimated distribution. Each month a house remains vacant on the rental market, with a certain probability its price is reduced as $r_s^t = r_s^{t-1} (1 - \exp(\varepsilon))$, where ε is a normal noise. Finally, the length of a rental agreement is chosen uniformly at random between 12 and 24 months. For details on all parameters and estimated distributions see Appendix B Table 12.

At the end of each tenancy agreement and for as long as the property remains unoccupied, the BTL owner of the house will consider whether to sell it or not. Similarly to the decision to buy new investment properties, the decision to sell those currently owned is also based on the computation of an investment yield V_{sell} . However, in this case the investment yield is computed for a specific instead of a hypothetical house, and thus for specific house quality and mortgage characteristics. It can therefore be considered as an effective yield for the property under consideration. In particular, the leverage can now be written as the current exponential average sale price of houses of the same quality, $\overline{p_Q}$, divided by the current equity stake of the household in the house, k , which can be computed in a mark-to-market style as the current exponential average sale price of houses of the same quality minus the remaining mortgage principal. Furthermore, an effective rental yield s can now be computed for this specific property by multiplying its current rental price by the expected average occupancy of houses of the same quality, and dividing this by the current exponential average sale price of houses of this quality. Finally, the mortgage cost can now be computed with the actual annual mortgage payment $12m$ divided by the current equity stake of the household in the house, k . As before, the effective yield can be written as a weighted sum of both rental and capital yield rates times the leverage minus the mortgage cost,

$$V_{\text{sell}} = \frac{\overline{p_Q}}{k} (\delta_i g + (1 - \delta_i) s) - \frac{12m}{k}, \quad (10)$$

where g is the expected annual house price growth [see Eq. (3)] and δ_i is the weight given by the specific investor type i to capital gains as opposed to rental yield. In this way, we can write the monthly probability of deciding to sell an empty investment property as one minus the monthly

probability of deciding to keep it,

$$P_{\text{sell}}^{\text{BTL}} = 1 - \sigma(\xi V_{\text{sell}})^{\frac{1}{12}}, \quad (11)$$

where ξ is a sensitivity parameter, setting how noisy or deterministic the decision is. If the investor decides to sell the property, she will place an offer in the sales market at the price given by Eq. (6).

3.4 Mortgage lending and its regulation: the bank and the Central Bank

The bank, which represents the mortgage lending sector in the aggregate, has the main role of extending mortgages to households. In doing so, it applies its own internal policy regarding underwriting standards, which reflect its internal risk appetite. In practical terms, these underwriting standards translate into constraints on the maximum amount of mortgage credit that the bank would be willing to extend to each potential borrower given their income or financial wealth. In particular, the bank considers four types of mortgage constraints, with parameters estimated based on the UK banking sector (see Appendix B Table 13 for estimation details):

- **Loan-to-value (LTV) limit.** A maximum value Γ_i is set for the ratio of the principal borrowed q over the value of the collateral, which is assumed to be equal to the transaction price of the house to be bought, with i representing the type of household (i.e., first-time buyer, home mover or buy-to-let investor). This sets a maximum value for the principal the household can borrow q as a function of the down-payment to be made d . Since the maximum possible down-payment is the household's full financial wealth w , the principal q must satisfy

$$q \leq \frac{\Gamma_i}{1 - \Gamma_i} w. \quad (12)$$

- **Loan-to-income (LTI) limit.** A maximum value Φ_i is set for the ratio of the principal borrowed q over the household's gross annual income y , with i representing the type of household. Thus, the principal to be borrowed q must satisfy

$$q \leq \Phi_i y. \quad (13)$$

- **Debt-service-to-income (DSTI) limit.** A maximum value Ψ is set for the ratio of the monthly mortgage payment over the household's gross monthly income $y/12$. This represents an affordability test. Writing the monthly mortgage payment as a function of the monthly interest rate $r/12$, the number of payments n and the mortgage principal q , the principal q must satisfy

$$q \leq \Psi y \frac{1 - (1 + r/12)^{-n}}{r} . \quad (14)$$

- **Interest-coverage-ratio (ICR) limit.** For BTL mortgages, a minimum value Ω is set for the ratio of the expected annual rental income over the annual interest expenses. This limit ensures that the rental income earned on the property exceeds the interest cost of the mortgage by a safety margin, which is meant to cover taxes, maintenance, changes in interest rates, or periods of vacancy, among others. Writing the expected annual rental income as the price of the house to be bought times the exponential moving average rental yield over all house qualities \bar{s} and bearing in mind that the maximum possible down-payment is the household's full financial wealth w , the principal q must satisfy

$$q \leq \frac{w}{\Omega \frac{r}{\bar{s}} - 1} . \quad (15)$$

Apart from these internal policy constraints, the mortgages extended by the bank must also comply with the lending regulation imposed by the Central Bank, which can also set LTV, LTI, DSTI and ICR limits. Each regulation set by the Central Bank can take two different forms: either it can be defined as a hard limit, preventing any lending beyond the corresponding hard cap; or it can be defined as a soft limit, allowing for a certain fraction of new mortgages to exceed the corresponding soft cap. This allowance is calculated over a 12-month rolling window. Note that, in the latter case, the bank will endogenously decide how many offers—mortgage in principle letters—over the soft limit to issue each month so as to remain under the maximum fraction once the market clears,¹⁵ and it will assign these high-LTI offers on a first-come-first-served basis.

¹⁵ Since it is impossible to know beforehand how many of the mortgage in principle letter holders will actually succeed in buying a house expensive enough to require a mortgage over the soft limit, there can be small and temporary breaches of the maximum fraction.

Depending on the type of borrower, the bank offers two different mortgage products:

- **Loans to owner-occupiers.** The only mortgage product that the bank offers to owner-occupier borrowers is a fixed-rate repayment mortgage with a maturity of 25 years,¹⁶ which is reduced depending on the borrower's age in such a way that all principal is repaid by retirement age (i.e. 65 years old). Mortgages to owner-occupiers have to comply with three constraints: (i) an LTV limit, (ii) an LTI limit, and (iii) a DSTI or affordability limit.¹⁷
- **Loans to BTL investors.** The only mortgage product that the bank offers to BTL borrowers is a fixed-rate interest-only mortgage with a fixed maturity of 25 years, which can only be offered until the borrower's retirement age (i.e. 65 years old). Note that a significant portion of BTL mortgages in the UK are interest-only and that the repayment of the principal is due only at the end of the term.¹⁸ Mortgages to BTL investors are subject to two constraints: (i) an LTV limit, and (ii) an ICR limit.

Importantly, there are no supply-side credit constraints in this model, e.g. there is no endogenous factor that could limit the supply of credit, such as liquidity or solvency constraints. Hence, there is no limit in principle regarding fiat money creation beyond the demand-side constraints already mentioned, and thus all credit applications complying with these borrower-based limits are always approved by the bank. The only endogenous lever for the bank to influence household demand for credit is the interest rate spread, which it sets in response to changes in the demand and with the goal of reaching a target level of credit per household (see Appendix A.5 for further details).¹⁹ In this way, the bank reacts to increases (decreases) in the

¹⁶ In the UK, it is unusual for mortgage products beyond a 10-year term to have interest rates that are fixed for the whole term (e.g. 25 years). However, as monetary policy is exogenous in our model and there is no default, flexible interest rates would only have a minor effect on housing market dynamics. A more realistic interest rate set-up is deferred to future research within a model with monetary policy and default.

¹⁷ In reality, the affordability test is often based on a stressed interest rate which is higher than the actual product rate, thereby leading to a more stringent test. The rationale behind a stressed rate is to account for changes in interest rates that may adversely affect a borrower's ability to make payments in the future if interest rates were to rise. Since the Central Bank policy rate is kept fixed and the bank spread fluctuates closely around its estimated value for our target calibration year, there is no need for such a stressed rate in our model.

¹⁸ Since BTL investors have no explicit saving motive for this final repayment of the principal, a mechanism is needed to prevent them from facing bankruptcy at the end of the term if they are not able to find a buyer for the property right away. In particular, from two years to maturity, they put the property up for sale whenever their financial wealth is not enough to repay the owed principal in full.

¹⁹ Note that, while the aforementioned underwriting standards do also influence the demand for credit, these

demand for credit with increases (decreases) in the interest rate spread of the mortgage products it offers. Finally, note that the interest rate paid by households on their mortgages is the sum of the described variable interest rate spread set by the bank and a fixed policy rate exogenously set by the Central Bank. That is, monetary policy is considered as exogenous and constant in our model.

3.5 Housing markets

The model contains two separate markets: a sales and a rental market. After households have sent their sale and rental bids and offers to the respective markets, these markets are cleared according to a double auction process, with the sales market being cleared first, followed by the rental market, so that BTL investors can directly offer their newly acquired investment properties for rent. The double auction process in each market proceeds in a number of rounds or iterations, where each iteration consists of two phases: *(i)* a matching phase, where each bid is matched to a specific offer; and *(ii)* a selection phase, where a specific bid is selected for each offer having received any bids.

In the sales market, owner-occupier bids are matched to the offer with the highest quality that they can afford and, within a given quality, to the cheapest house available. BTL bids, on the contrary, are matched to the offer with the highest expected gross rental yield that they can afford. In the rental market, all bids are matched to the offer with the highest quality that they can afford and, again, cheaper houses are preferred in case of equal quality. As a result, in both markets, some offers are matched to one or multiple bids while others might not be matched to any. Similarly, bids with too low a bidding price are not matched to any offer.²⁰

In the selection phase of both markets, the procedure iterates through the offers that were matched to any bids. For offers that were matched to a single bidder, that bidder is selected and the transaction is completed at the price demanded by the seller. For offers that were matched to multiple bidders, the seller has a certain probability to increase or bid up the price by a small

are exogenous fixed limits and thus cannot be used to address endogenously emerging temporary over- and undershootings of the credit supply target.

²⁰ In terms of the informational structure of the markets, these rules imply that the matching phase is characterised by perfect information, in the sense that bidders have access to information about the price and quality of all available offers, such that they can find their optimal bidding option given their bidding price.

amount (see Appendix A.6 for details). A bidder is then randomly chosen among those who can still afford the offer and the transaction is completed at the updated offer price.²¹

Unsuccessful bids and offers are returned to their respective market general pool, and the process continues in more matching and selection rounds until (i) there are no more bids left, (ii) there are no more offers left, or (iii) the bids that remain cannot be matched with any of the offers that remain, their bid prices being below the available offer prices. At the end of the clearing process, all remaining unsuccessful bids are deleted, that is, the households who made them will have to bid again in the following month if they so decide. On the contrary, all remaining unsuccessful offers are kept for the following month, although the households who made them might decide later on to update their offer price or remove the offer from the market.


4 Validation

Significant advances have been made in recent years regarding the validation of agent-based models (ABMs) (Fagiolo et al., 2019). In this paper, we assess how the output of the model matches with key features of the housing and mortgage market observed in the data for the United Kingdom. Our comparisons comprise broad patterns and stylised facts (such as the emergence of house price cycles), but also quantitative measures, such as the mean and range of various housing and mortgage market indicators. This broadly follows the approach taken by the macro ABM literature, for example Dawid et al. (2012), Ashraf et al. (2017), and Popoyan et al. (2017). One difference is that the literature mostly seeks to match first and second moments, while our focus is also on distributions. Distributions are important in the context of the policy experiments below, in particular the soft loan-to-income flow limit, in which a certain share of borrowers is not affected by the policy. Due to space constraints, we only show some of these validation checks here.

One of the key features of the model is its ability to generate real-world-like house price

²¹ We choose the winning bid at random, and not simply select the bidder with the highest bidding price, to reflect the fact that the maximum bidding price (i.e. the maximum price a bidder is willing to pay) is considered private information, and thus it is not known to the seller. Note that this random selection plays the role of a market imperfection, in the sense that households with a higher maximum bidding price (usually wealthier) may occasionally not be selected among multiple bidders and could end up having to bid for a lower quality house (usually cheaper).

cycles which emerge naturally from the interactions of buyers and sellers in the market. Figure 1 shows the house price index in a sample simulation run of 50 years. The model appears to match the amplitude and frequency of the observed data reasonably well. While we do not incorporate in the model all factors that affect house price cycles in reality, such as credit supply conditions, we are still able to broadly match the pattern found in the data.



Source: OECD.Stat database. United Kingdom, real house price index, seasonally adjusted. Data are detrended with the help of a Hodrick-Prescott filter, using a smoothing parameter, λ , of 129600 (360^2). An equivalent time period is taken from simulated data of a typical run of the model.

Figure 1: House price cycles

While household expectations, encapsulated in the house price growth expectation in Equation (3), are the main factor behind the house price cycles in the model, heterogeneity is also fundamental for cycles to appear in this framework.²² In particular, income and wealth heterogeneity, as well as the existence of an active buy-to-let sector, play an essential role in the emergence of strong cycles, with realistic peak-to-trough amplitudes. The specific role of these factors, however, will vary across the different phases of the cycle. For instance, at the beginning of the boom phase, as prices start increasing and, with them, also expected capital gains, more and more households are driven to bid in the ownership market. As prices continue to rise and more constrained households (lower income, lower financial wealth) are progressively priced out of the market, less constrained households (higher income, higher financial wealth) are still able to bid and absorb the available houses on offer, thereby driving further price in-

²² In fact, if either expectations or household heterogeneity are deactivated (respectively, by setting the expected price growth to zero and by assigning all households the same income, propensity to save and investor type), the cycles completely disappear and the house price index behaves as small fluctuations around the mean.

creases. Moreover, price growth drives capital-gains-seeking buy-to-let investors to keep their properties and even try to acquire more, thereby at the same time increasing the demand and reducing the offer. In this way, household heterogeneity allows for longer and stronger boom phases, as demand moves from being sustained by the whole household sector to being more and more concentrated among wealthier households and buy-to-let investors. As prices increase further, at some point, the simulated housing market starts to run out of wealthy enough households able and willing to buy new properties. As a consequence, the demand becomes unable to absorb the available houses on offer, and so the price trend reverses and prices start falling, initiating the bust phase of the cycle. As prices drop and expected capital losses soar, renting becomes more and more attractive with respect to buying, which keeps the demand down even if prices are progressively becoming more affordable. At the same time, expected capital losses also drive capital-gains-seeking buy-to-let investors to decrease their portfolios of properties, thus both reducing the demand and increasing the offer in the ownership market. Finally, as prices decrease further, more and more households become progressively unconstrained and could potentially afford to bid in the ownership market. Even if their probability to do so is very low due to the expected capital losses, as the number of potential buyers increases sufficiently, at least some of them will decide to buy. This will gradually moderate, eventually stop and finally reverse the price trend, thereby triggering a new cycle. Again, income and wealth heterogeneity are crucial for this process to be sufficiently gradual, thereby allowing for larger price drops and longer bust phases.²³

We further compare the simulation results with several empirical housing and mortgage market statistics. We use data from 2011 as our baseline mainly for three reasons: first, the year 2011 precedes the housing market policies (a loan-to-income policy and an affordability test) introduced by the Bank of England in 2014. As we perform similar policy experiments in Section 5, this allows us to more clearly assess their effects compared with an environment with no such policies in place. Second, 2011 represents a point in time that can be considered

²³ Note that, if all households had the exact same (or very similar) income and wealth, then they would all become constrained or unconstrained at the exact same (or very similar) price level. As a consequence, any increasing price trend above that level would be quickly met by a drastic drop in demand, while any decreasing price trend below that level would be quickly counteracted by a strong rise in demand, thereby not allowing these initial trends to develop into realistic boom and bust episodes.

as comparatively 'neutral' from a housing cycle perspective in the UK, i.e. neither a boom nor a bust episode. This allows for a relatively natural benchmark for our model. Third, 2011 coincides with the best data availability between the Great Recession and the introduction of the Bank of England's housing market policies, in particular in relation with certain survey data.²⁴

The UK housing market can be characterised by a significant gap between the rate of house building and growth in demand, and growth in house prices in excess of earnings. In 2011, around 65 percent of the housing stock was owner-occupied,²⁵ and the share of property wealth (excluding private pension wealth) of total wealth was 60 percent.²⁶ UK household mortgage debt has been historically high relative to income, in particular, equal to 103.8 percent of household income by 2011 (see Table 1).

Table 1 provides a comparison of average monthly house prices, the number of housing transactions and mortgage approvals, and the mortgage debt-to-income ratio generated by the model (averaged over 50 Monte Carlo simulations) and real data in 2011. Apart from the number of housing transactions, the other three indicators are very close to the 2011 values. Note that the average number of housing transactions is well within the range of the minimum and maximum values observed in the data between 2005 and 2014.²⁷

	Simulation	2011	2005-2014	
	Mean	Mean	Min	Max
(a) average house prices (£1,000)	171.8	167.9	153.9	193.2
(b) housing transactions (1,000)	100	73.7	51.6	149.4
(c) mortgage approvals (1,000)	51.9	49.3	26.6	129.1
(d) mortgage debt-to-income ratio (%)	98	103.8	94.3	109.4

Source: (a) UK average house prices are obtained from the Office for National Statistics (ONS). (b) Housing transactions are the number of residential property transactions in the United Kingdom with a value of £40,000 or above per month (HM Revenue and Customs). (c) Seasonally adjusted mortgage approvals for sterling loans secured on dwellings, net of cancellations (Bank of England). For (b) and (c) see underlying data of Bank of England June 2014 Financial Stability Report Chart 2.8. (d) Mortgage debt-to-income ratio (Bank of England), see underlying data of Bank of England November 2017 Financial Stability Report Chart A.9.

Table 1: Housing and mortgage market indicators

²⁴ These are also the reasons why we calibrate the model with 2011 data, where available. See Appendix B and the Online Appendix for further detail on the data used for the calibration.

²⁵ Source is DCLG Table 101 - live tables on dwelling stock (including vacants).

²⁶ Source is ONS Wealth and Assets Survey - Wealth in Great Britain, Table 2.1.

²⁷ The choice of the 2005-2014 period is based on several factors: (i) data are available from 2005 on; (ii) it covers a full boom and bust cycle; and (iii) it predates the impact of the Bank of England's loan-to-income and affordability policies introduced in 2014.

Table 2 shows that the average loan-to-value (LTV), loan-to-income (LTI), and house-price-to-income ratios as well as the age²⁸ of owner-occupier mortgage borrowers generated by the model (averaged over 50 Monte Carlo simulations) are in line with the averages of these variables obtained from a comprehensive loan-level mortgage dataset. Figure 2 (a-d) shows the distributions of the same variables, which further supports a good qualitative fit between the model and the real data.

	Simulation	PSD
Mean LTV ratio (%)	65.8	68.3
Mean LTI ratio	2.84	2.97
Mean house-price-to-income ratio	4.4	4.7
Mean age (years)	35.3	35.7

Source: The Financial Conduct Authority’s (FCA) loan-level Product Sales Data (PSD), which includes regulated mortgage contracts only. The above statistics are calculated with 2011 data.

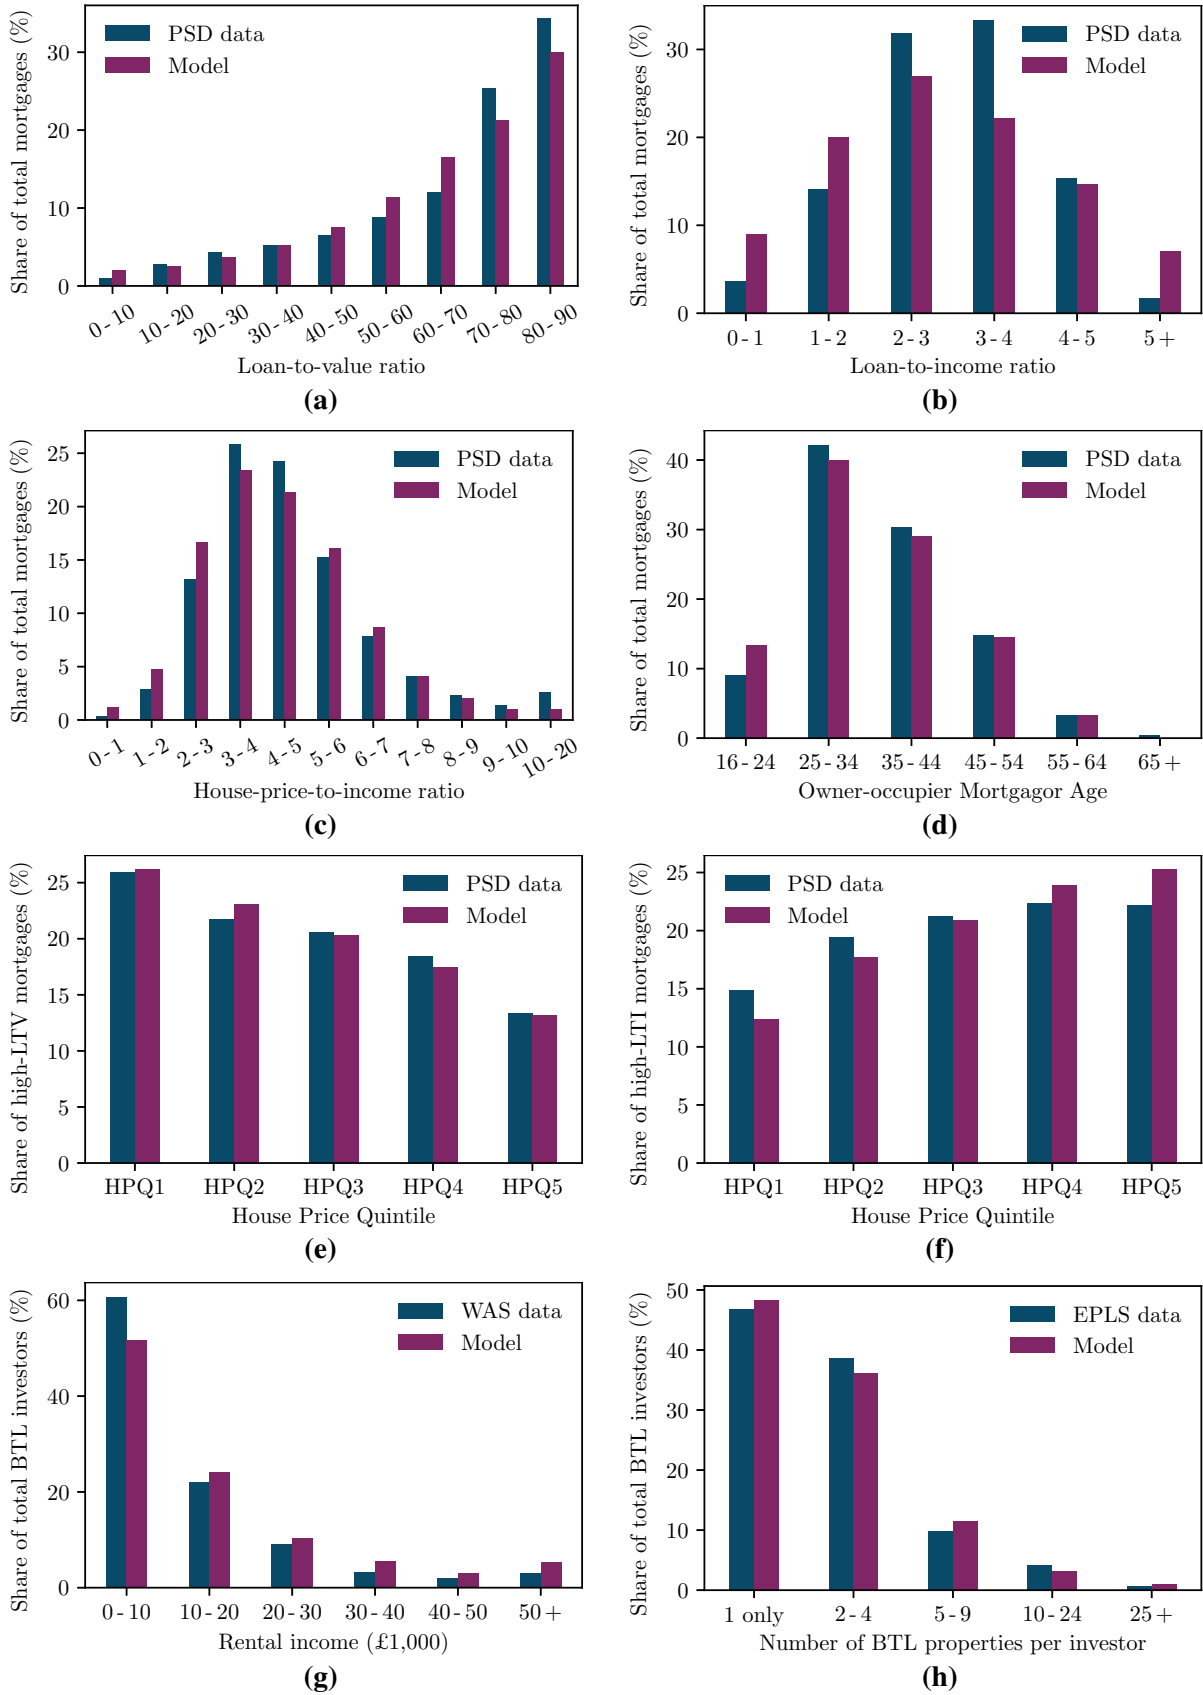
Table 2: Owner-occupier mortgage characteristics

We also compare our simulation results against real data to understand the composition of risky lending, i.e. loans with a high LTV and/or LTI ratio²⁹. First, we group owner-occupier house transactions (financed by mortgages) into five house price quintiles. Figure 2 (e) shows, for high-LTV lending, the share of mortgages by house price quintile, e.g. 26% of high-LTV loans finance houses in the lowest quintile. Similarly, Figure 2 (f) shows the share of mortgages by house price quintile for high-LTI lending, e.g. 22% of high-LTI loans finance houses in the highest quintile. Our simulation results are consistent with real data both in terms of the shares in each house price quintile as well as the monotonic relationship between the lowest and highest quintile.

A realistic rental market is one of the key features of our model. We provide a comparison between simulation and real data for rental income and the distribution of buy-to-let (BTL) properties. Figure 2 (g-h) shows that our simulation is comparable with real data regarding the distribution of rental income earned by BTL investors and the number of BTL properties held per BTL investor.

²⁸ Note that while data on households’ age distribution is used to calibrate the model, the mortgage borrowers’ age distribution is endogenous due to their housing and mortgage decisions.

²⁹ High-LTV lending is defined as an LTV ratio of 75% and above. High-LTI lending is defined as an LTI ratio of above 3.35 - the threshold for our policy experiment as described in Section 5 - to provide a consistent comparison.



Source: (a)-(f) The Financial Conduct Authority's (FCA) loan-level Product Sales Data (PSD, 2011), which include regulated mortgage contracts only; (g) Wealth and Asset Survey (WAS) Wave 3 (2010-2012); (h) English Private Landlord Survey (EPLS, 2018). All model results are averages over 50 Monte Carlo simulations.

Figure 2: Simulation vs real data - Comparison of key mortgage and rental market variables

5 Experiments

We use the model to perform a series of comparative statics exercises to investigate the impact of macroprudential housing policies on the mortgage, housing and rental markets. These are hypothetical experiments that do not represent any actual policies by the Bank of England. We conduct two policy experiments by imposing regulatory restrictions on households' ability to access mortgages: (i) a hard loan-to-value (LTV) limit, and (ii) a soft loan-to-income (LTI) limit.

In the benchmark case, the maximum LTV ratio that the bank applies to owner-occupier mortgages is 90 percent. In the first experiment, the maximum LTV ratio for owner-occupier mortgages is capped at 85 percent by the Central Bank. In the second experiment, the Central Bank imposes an LTI soft limit of 3.35,³⁰ but allows the bank to issue 15 percent of new mortgage lending with an LTI above this limit.³¹ The allowance is calculated over a 12-month rolling window. Note that the bank also maintains its hard internal LTI limit (i.e. a maximum LTI of 5.6 for home movers and of 5.4 for first-time buyers).

In the following subsections, we analyse how a range of mortgage, housing and rental market indicators are affected by these experiments, compared with the benchmark of no policy. We consider how the policy experiments affect these variables in the aggregate, how they change their time-series behaviour, and how they affect different types of households. Note that all results presented within this Section correspond to averages over 50 Monte Carlo simulations.

³⁰ This limit is chosen in a such a way that the percentage of borrowers affected by the LTI policy experiment will be similar to the LTV policy experiment, e.g. in the benchmark case around 20 percent of borrowers have mortgages with LTV 85 and above. To map this into the LTI policy experiment we choose a limit above which there are 20 percent of borrowers (after deducting the flow allowance).

³¹ The Bank of England introduced an LTI policy of 4.5 with an allowance of 15% in 2014 (see e.g. [Bank of England, 2019](#)). This policy was not calibrated to reduce mortgage lending at the time when it was introduced. Instead, its aim was to provide insurance against the risk of a loosening in underwriting standards and an increase in household indebtedness in the future. In contrast, we calibrate the policy experiment in this paper to have a marked effect on the housing market to demonstrate the capabilities of our model. We therefore do not seek to assess the effects of the Bank of England's LTI policy.

5.1 Aggregate effects

Table 3 presents the aggregate effects of the policy experiments³² on several key variables. As expected, the LTV and LTI policy experiments lead to a significant decrease in the average LTV and LTI ratios, respectively, as the policies directly impact the availability and size of these loans to borrowers (we will look at the LTV and LTI distributions in more detail in Section 5.3).

	Benchmark	LTV cap	LTI cap
LTV ratio (%)	65.8	63.6	63.3
LTI ratio	2.84	2.71	2.4
# of mortgage approvals	51,898	50,085	52,100
Debt-to-income ratio (%)	98	89.1	87.9
# of housing transactions	100,006	100,433	100,500
Sale price (£)	171,768	166,473	164,655

Note: These statistics are averages over 50 Monte Carlo simulations.

Table 3: The aggregate effects of policy experiments on key indicators

Compared to the benchmark case, mortgage approvals fall by an average of 3.5% under the LTV policy experiment, while the LTI policy experiment does not make a significant difference. One explanation is that the LTV policy is a hard limit, while the LTI policy allows for a 15% share of mortgages to be issued above the set limit. The LTV policy considered here may therefore exclude some borrowers permanently from the mortgage market, while the bank has more flexibility to accommodate borrowers under the soft LTI policy. Interestingly, the aggregate debt-to-income ratio falls by about the same amount, 10 percentage points, under both policy experiments. This, together with the differences in mortgage approvals, hints at the fact that the size of the average mortgage the bank grants is higher under the LTV than the LTI policy. The policy experiments may therefore have distributional consequences at the extensive and intensive margin, i.e. some households may not be able to receive mortgages, and the size of mortgages available to them may change.

The number of housing transactions proves remarkably stable in the presence of the policy experiments. The reason is that, in our model, housing transactions are essentially determined by structural factors over the long run, such as demographics (which we do not vary during

³² By aggregate effects we mean differences in the long-term averages of the relevant variables. The averages can be interpreted as the long-run equilibrium values around which the variables fluctuate. Therefore, the values shown should not be seen as a short-run impact of the introduction of the policies.

the simulation). Additionally, cash buyers and buy-to-let (BTL) investors who are not affected by these policy experiments can fill the gap in the housing market. We assess this further in Section 5.3 and Section 5.4.

Average sale prices are lower under both policy experiments, falling by about 3.1% for the LTV and 4.1% for the LTI limit compared with the benchmark case. This is consistent with the literature, which tends to find that an increase in the availability of credit leads to higher house prices, especially if housing supply is inelastic, as in our model (see for example [Adelino et al., 2012](#); [Favara & Imbs, 2015](#)), so we would expect the introduction of borrowing caps to reduce credit supply and house prices.

5.2 Housing booms and busts

While the previous section investigated changes to the long-run averages of several model variables due to the borrowing limits, this section considers how these variables are influenced by the policy experiments depending on the house price cycle. This matters as researchers are not only interested in how policies affect the averages of variables of interest, but also whether they can mitigate volatility. We first identify housing boom and bust periods in the model by computing the trend component of the house price index using a Hodrick-Prescott filter.³³ We determine where this trend component is increasing or decreasing to identify boom and bust episodes. We then condition the variables we investigate below on these episodes.

We focus on three variables: mortgage approvals, the number of housing transactions, and house prices. The policy experiments mitigate fluctuations in these variables across the house price cycle, i.e. they tend to be lower during booms but higher during busts. Figure 3 expresses this as the percentage deviation from the average over the cycle for a given experiment. Under the benchmark of no policy, mortgage approvals are higher by about 34% in booms and lower by 34% in busts.³⁴ But when the LTV and LTI policies are in place, these numbers fall to 26% in booms and -26% in busts, respectively. Housing transactions follow a similar pattern: with the policies in place, they increase by about 5 to 6 percentage points less during housing booms

³³ We use a smoothing parameter, λ , of 100,000.

³⁴ All numbers behave broadly symmetrically in booms and busts, though it should be kept in mind that the long-term averages are different for each policy experiment, in line with the results in Table 3.

and also decrease less by the same margin during housing busts compared with the benchmark of no policy. So, while the average number of housing transactions stays about constant in the long term as discussed in the previous section, the timing of when these transactions take place is influenced by constraints on the availability of credit depending on the state of the housing cycle. Last, house prices also behave in a similar way, but to a lesser degree: when the policies are in place, they increase by about 2.5 percentage points less during booms and decrease by about 2.5 percentage points less during busts.

We can also confirm a finding by [Tarne et al. \(2021\)](#). When looking at the level of house prices (see Table 4), the policies are most binding when house prices are exuberant, i.e. during booms: the LTV limit and the LTI limit lead to a drop of 5.1% and 6.1%, respectively, in average house prices, compared with the benchmark of no policy. But the policy experiments have less of an effect when house prices are relatively low, i.e. during busts: they only lead to a drop of 0.3% and 1.4%, respectively, in average house prices. That is, most of the effect of the policy experiments on house prices takes place during housing booms.




Figure 3: Key variables in housing booms and busts

	Aggregate	Boom	Bust
Benchmark (£1,000)	171.8	201.5	141.7
LTV cap (£1,000)	166.5	191.1	141.3
Change w.r.t. benchmark	-3.1%	-5.1%	-0.3%
LTI cap (£1,000)	164.7	189.1	139.7
Change w.r.t. benchmark	-4.1%	-6.1%	-1.4%

Note: These statistics are averages over 50 Monte Carlo simulations.

Table 4: Effect of policy experiments on house prices

To corroborate our finding that the policy experiments mitigate fluctuation across the cy-

cle, we consider two additional variables: credit growth and house price growth (not shown). Credit growth is higher during house price booms and lower during house price bust, consistent with the literature (e.g. [Jordà et al., 2016](#)). The policy experiments have a mitigating effect on this credit cycle: the average credit growth rate falls from about 3% to 2.3% for LTV and 2.4% for LTI during housing booms and increases from -2.9% to -2.3% for both LTV and LTI during housing busts. House price growth follows a similar pattern as credit growth. As expected, it is lower (negative) in busts and higher (positive) in booms. Hence, these macroprudential policy experiments reduce fluctuations in credit growth and house price growth over the cycle.³⁵

In summary, the LTV and LTI policy experiments reduce the availability of mortgage loans, which is reflected in lower debt ratios and house prices. But they also reduce fluctuations in many of these variables across the house price cycle. While in the aggregate the policy experiments have little influence on variables that are determined by structural factors such as the number of housing transactions, they may still have the desired effect of reducing the amplitude of fluctuations over the cycle.³⁶

5.3 Cross-sectional effects

In addition to analysing the effects of policy experiments in mitigating the cyclical behaviour of economic variables, it is important to understand how policy measures affect different types of households from a cross-sectional perspective. This is because policy experiments might have disproportionate or unintended side effects on certain groups of households.

LTV and LTI distributions. We begin by investigating the average of the LTV distribution. As shown in [Table 5](#), the average LTV ratio for owner-occupier mortgages is around 2pp lower when the LTV and the LTI caps are in place. The adjustment is larger for first-time buyers (FTBs) than for home movers (HMs). This is to be expected as FTBs are more bound by the

³⁵ As an alternative metric, we can calculate how the policy experiments affect the standard deviation of credit growth and house price growth. In the case of credit growth, it falls from about 4.2% to 3.3% for the LTV cap and 3.5% for the LTI cap. In the case of house price growth, it declines from 3.77% to 3.36% for the LTV limit and 3.56% for the LTI limit.

³⁶ There is no clear mitigation of the cyclical pattern for the LTV and LTI ratios. We only find a reduction of the cycle for the LTI ratio under the LTI policy experiment. One explanation is that many variables, such as loan amounts and house prices, are affected by the policy experiments to a similar degree, so the ratios of such variables tend to remain relatively stable. One exception is income, which is exogenous in our model and could therefore explain the behaviour of the LTI ratio under the LTI policy experiment.

policy experiments, especially in terms of the LTV limit.³⁷ Under the LTI cap, the average LTI ratio decreases by 4 pp, affecting average FTB and HM LTI ratios similarly.

	Benchmark	LTV cap	LTI cap
Average LTV (%)			
Owner-occupier	65.8	63.6	63.3
FTB	68	64.4	65.2
HM	62.8	62.5	60.7
Average LTI			
Owner-occupier	2.8	2.7	2.4
FTB	2.7	2.4	2.2
HM	3.1	3.1	2.7

Note: These statistics are averages over 50 Monte Carlo simulations.

Table 5: Average LTV and LTI ratios by borrower types

By virtue of the heterogeneity embodied in agent-based models, we can perform a more granular investigation across the whole distribution of the LTV and LTI ratios. This is particularly relevant as financial risks are often disproportionately concentrated in high-LTV and high-LTI loans, i.e. in the tail of the respective distributions.³⁸ In other words, focusing on the average LTV and LTI ratio to assess financial stability risks may give an incomplete picture of the true riskiness of the bank’s mortgage portfolio and of the effects that policy measures may have.

Figure 4 shows the upper end of the LTV distribution, the “tail” (for LTV ratios of 75-80%, 80-85%, and above 85%), in the benchmark case and when the LTV policy is in effect. The share of high-LTV mortgages for FTBs (with an LTV ratio above 75%) decreases from 52 percent to 45 percent. The share of mortgages with an LTV ratio between 80 and 85% increases significantly and accounts for two thirds of all high-LTV lending above 75% when the LTV policy is in place. For HMs, the share of high-LTV lending (above 75%) is about the same as in the benchmark case, with the 75-80% and 80-85% LTV bands absorbing the capped 85+ LTV

³⁷ Wong et al. (2011) find that mortgage insurance policies (MIPs) can mitigate the impact of an LTV policy on borrowers. MIPs allow mortgage lenders to issue mortgages with LTV ratios above the maximum threshold. These policies aim to protect lenders from credit losses on the loans above the maximum LTV threshold in case of default and transfer the risk from lenders to insurers. In the UK, such MIPs are not currently available, therefore we do not investigate this option.

³⁸ For example, Aron & Muellbauer (2016) assess the proportion of mortgages in negative equity, i.e. with an LTV ratio above 100%, as a determinant of reposessions and arrears.

share about proportionally.

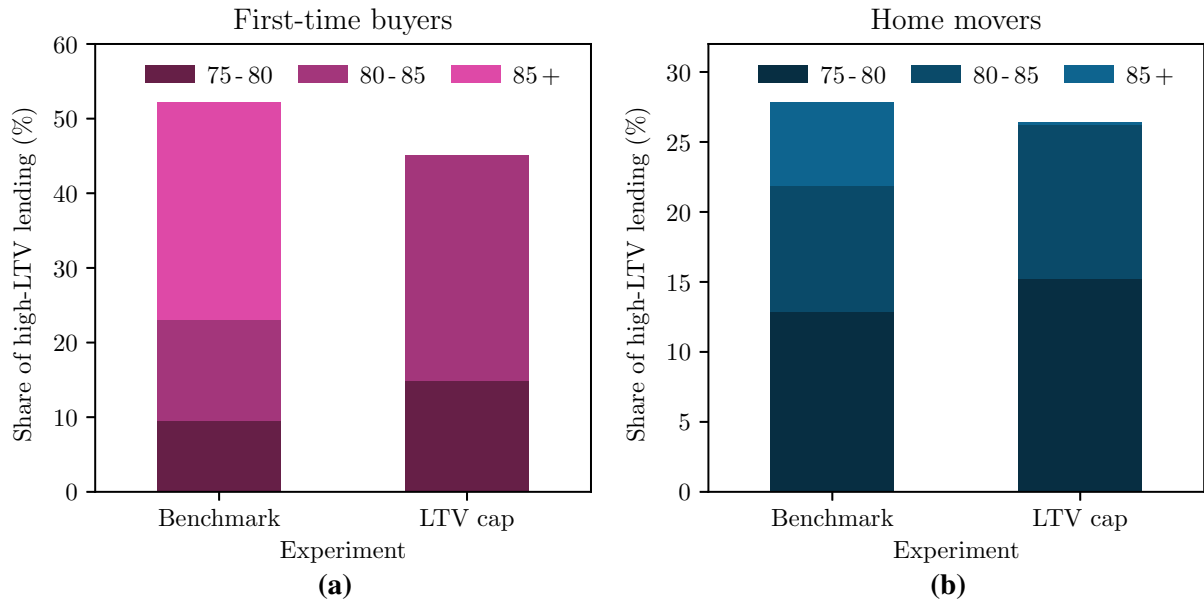


Figure 4: Share of high-LTV lending by borrower type

Similarly, Figure 5 shows the tail of the LTI distribution (for LTI ratios of 3.0-3.35, 3.35-4.0 and above 4.0). It is important to recall that the LTI policy experiment is a soft limit, so there are mortgages with an LTI ratio above the policy limit of 3.35, while there are no mortgages above the hard LTV cap of 85%. Several findings stand out. First, the impact of the LTI cap is stronger for FTBs than HMs as the share of high-LTI lending decreases by 22pp for FTBs compared to a 16pp reduction for HMs.³⁹ Second, the bank allocates a higher share of high-LTI loans to HMs than FTBs when the policy is in place as the share of HMs in high-LTI lending increases from 50% to 70%. Third, there is clustering right below the soft limit of 3.35 as many aspiring FTBs and HMs are pushed just below the limit. A similar bunching effect can be observed for the UK's LTI policy introduced in 2014 ([Bank of England, 2019](#)).

We go one step further by analysing how LTI and LTV ratios are jointly affected by housing policy experiments, not just in isolation. Figure 6 shows LTV distributions for FTBs and HMs. Each chart presents LTV distributions for the benchmark case and when the LTI cap is in place. Introducing an LTI cap also shifts the LTV distribution to the left for both FTBs and HMs, with the effect somewhat stronger for the latter. The share of high-LTV mortgages (i.e. LTV ratio above 75%) decreases from 52 percent to 46 percent for FTBs and from 28 percent

³⁹ The share of high-LTI lending (an LTI ratio greater than or equal to 3.35) for FTBs decreases from 30.3 to 7.9; and for HMs decreases from 42.3 to 25.9 (i.e. the top two segments of the stacked bars in each graph).




Figure 5: Share of high-LTI lending by borrower type

to 22 percent for HMs. This demonstrates an important lesson for the calibration of policies: a given policy does not only act on the risk metric it is directly targeting (such as the LTI ratio in the case of an LTI policy), but it may also affect other risk characteristics, such as high-LTV lending, as a side effect. Therefore, for a given desired reduction of risk, a holistic view of the joint distribution of risk characteristics is needed to ensure an appropriate calibration of an individual policy.




Figure 6: LTV distributions when LTI cap is in place

Housing Transactions. The number of housing transactions in the aggregate does not change significantly under the different policy experiments (Table 3). Figure 7, however, indi-

cates a compositional change as the share of transactions by buyer types differs across policy experiments. With the LTV (LTI) policy, the share of FTBs (HMs) decreases, and the share of BTL investors' house purchases increases. We also see a similar pattern with mortgage approvals,⁴⁰ indicating a spillover effect of owner-occupier housing policy experiments to the BTL mortgage market.




Figure 7: Change in housing sales transactions

We further investigate compositional changes in risky lending, in particular whether housing policy experiments lead to changes in the proportion of cheap vs expensive houses funded by high-LTV and high-LTI mortgages. This is relevant from a financial stability perspective, as the exposure of the bank to a risky loan (for a given LTV and LTI ratio) increases with the value of the house. To understand the compositional changes in risky lending and whether certain housing transactions are affected disproportionately by housing policy experiments, we first group the pool of all owner-occupier house transactions financed by mortgages into five house price quintiles, i.e. from cheapest (HPQ1) to most expensive (HPQ5) houses. Then we calculate, for each quintile, the number of high-LTV (75% and above), respectively high-LTI (above 3.35), transactions divided by the total number of high-LTV, respectively high-LTI, transactions

⁴⁰ Regarding mortgage approvals, the LTV cap hits mostly FTBs relative to HMs, and there is a significant increase in mortgage approvals to BTL investors. With the LTI soft limit the mortgage approval numbers are slightly lower for both FTBs and HMs, however advances to BTL investors increase by about 5 percent. This suggests that even a small decrease in owner-occupier mortgages could have a significant knock-on effect on the BTL mortgage market, which is a relatively small segment of the overall market.

across all quintiles.⁴¹

Figure 8 shows that in the benchmark case 25% of high-LTV (LTI) lending finances the purchase of houses in the lowest (highest) house price quintile. While the share of high-LTV lending decreases monotonically between the lowest and the highest house price quintiles, the share of high-LTI lending increases monotonically. In the high-LTV lending, FTBs constitute a higher share compared to HMs in all house price quintiles. On the other hand, the share of HMs increases over house price quintiles for high-LTI lending.




Figure 8: Share of high-LTV and high-LTI lending by house price quintile and borrower type

Figure 9 shows how this share differs, for each house price quintile, when the LTV/LTI policy is in place compared with the benchmark of no policy. The green bars represent all owner-occupier borrowers, while the purple and blue bars show the split by FTBs and HMs, respectively, which allows us to investigate the compositional changes in risky lending.⁴² The figure reveals that FTBs lose some of their access to high-risk mortgages when the policies are in place, while HMs increase their access. The extent of this loss of access depends on house prices: the more expensive a property becomes, the less likely FTBs are to secure a high-risk mortgage. The reason is that FTBs are more constrained than HMs in terms of their income and savings, so when policy limits tighten, FTBs are more likely to be affected. The only time

⁴¹ The formula used to calculate these shares is: the number of mortgages (respectively, housing transactions) in a given house price quintile financed by high-LTV (LTI) lending divided by the total number of high-LTV (LTI) mortgages (respectively, housing transactions). See Figure 2 (e-f) for a comparison of the benchmark case with real data.

⁴² Note that for each house price quintile the purple and blue bars add up to the green bar. Also across HPQ1-HPQ5, the green bars add up the zero.

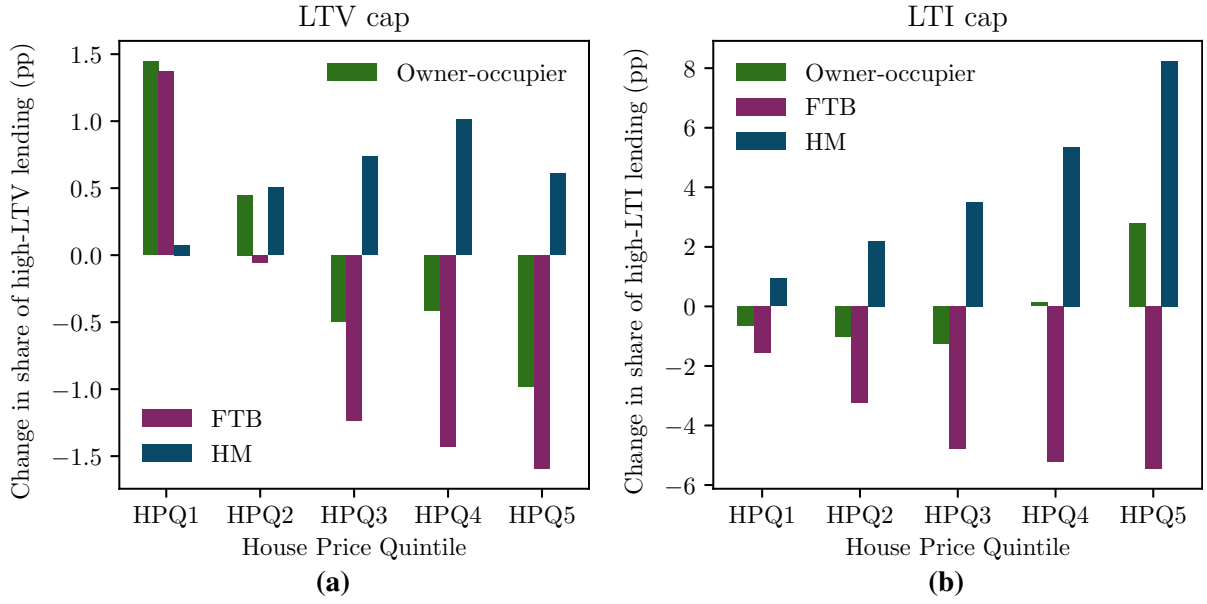


Figure 9: Change in the share of high-LTV and high-LTI lending by house price quintile

when FTBs can increase their share of high-risk loans is in the lowest house price quintile under the LTV policy, where they are least likely to be constrained relative to HMs. In the aggregate (green bars), the share of risky lending decreases the higher house prices are under the LTV policy, but increases under the LTI policy. This is because FTBs are more prevalent in the high-LTV market, while HMs dominate the high-LTI market in high house price quintiles (Figure 8).

In summary, housing policy experiments affect certain types of borrowers and house buyers disproportionately, with FTBs generally being affected more. We confirm that there is bunching just below the LTI policy limit. The policy experiments affect various risk metrics simultaneously, not just the one they are targeting directly. The policy experiments cause a shift of housing transactions away from owner-occupiers to the BTL sector. And the more expensive the property, the more HMs have access to the available pool of high-risk mortgages relative to FTBs.

5.4 Spillovers to the BTL sector and rental market

In the previous section, we showed that both housing policy experiments lead to a decline in owner-occupier mortgage approvals/housing transactions and to an increase in BTL mortgage approvals/housing transactions. The results are more prominent when the LTV cap is in place.

A decrease (increase) in FTB (BTL) purchases indicates that the housing policy experiments, in particular the LTV limit, could also have an impact on the rental market as the house ownership structure changes.

In the benchmark case, the share of owner-occupiers (renters) is 54 (38) percent. Figure 10 (a) shows that the share of owner-occupiers (renters) is lower (higher) when housing policies are in place. The LTV policy experiment leads to a 1.6 percentage points decrease in the ownership rate, moving the equivalent of 422 thousand households (out of 26 million) from home ownership to the rental market.

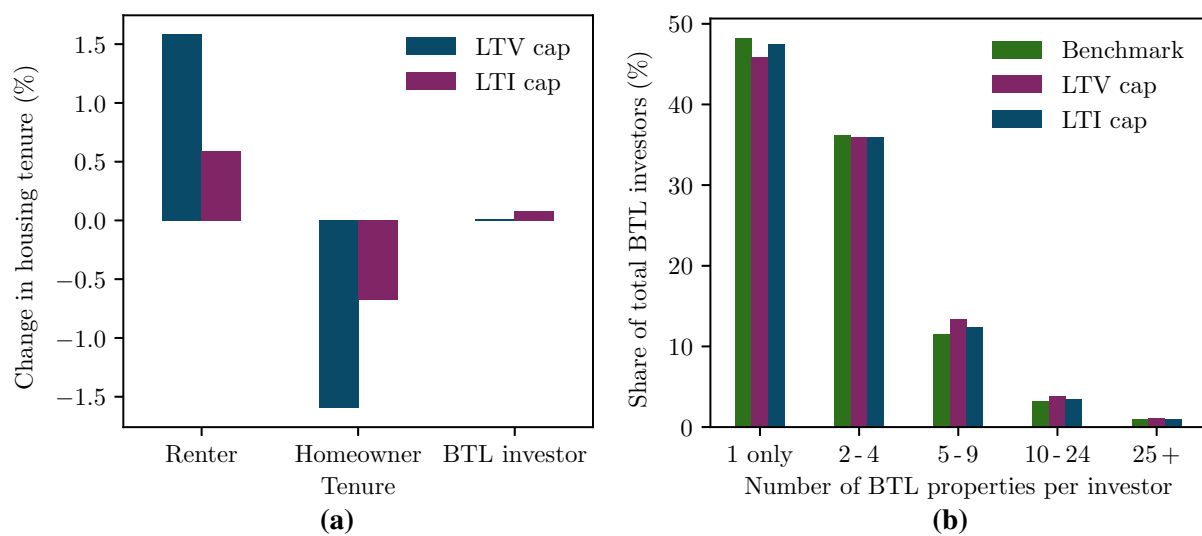


Figure 10: Change in tenure and distribution of the number of BTL properties

Changing our perspective from households to houses, the share of the housing stock held by BTL investors is higher by 2 pp when the LTV cap is in place (not shown). Given that the share of BTL investors does not change (Figure 10 (a)), this indicates that the distribution of the number of properties owned by BTL investors might change when housing policies are in place. Figure 10 (b) shows that the share of BTL investors who have multiple BTL properties, in particular between 5 and 9 properties, is higher when the LTV cap is in place. This suggests that BTL investors with bigger portfolios may benefit from the introduction of housing policies which only target owner-occupiers. The impact of the LTI policy experiment on BTL investors' portfolios is qualitatively similar to the LTV policy experiment, but less pronounced.⁴³

⁴³ Additionally, we look into the effect of housing policy experiments on housing wealth. The share of housing wealth owned by the top 10 percent of households increases under the LTV and LTI policy experiments by 0.71 pp and 0.77 pp, respectively. The share of housing wealth owned by the bottom 50 percent of households decreases by 0.33 pp under the LTV policy experiment and by 0.15 pp under the LTI policy experiment.

Policy experiments on owner-occupier mortgages limit the number of FTBs who can get on the housing ladder, therefore leading to an increase in rental demand. At the same time, these housing policy experiments also increase the rental supply as BTL investors increase their housing portfolio. A priori, this should increase the number of rental transactions, while the effect on prices (i.e. rents) is ambivalent. Indeed, Table 6 shows that the number of rental transactions increases. The average monthly rent increases when the LTV cap is in place as, based on our model calibration, the increase in demand dominates the higher supply. However, in case of the LTI policy experiment, the increase in supply and demand is more balanced, leading to a smaller increase in rental transactions and a very small decrease in rents. The rental yields are higher with both policy experiments due to significant decreases in house prices. Interestingly, the rent-to-income ratio is slightly lower with the policies in place. This is because the marginal tenant, who would be a FTB without a housing policy but ends up renting with a policy, has a relatively higher income than the other tenants who would be renters even without a housing policy in place.

	Benchmark	LTV cap	LTI cap
# of rental transactions	327,662	351,564	337,221
Monthly rent (£)	526	534	523
Rent-to-income ratio (%)	25.1	24.8	24.8
Rental yield (%)	4.2	4.5	4.3

Note: These statistics are averages over 50 Monte Carlo simulations.

Table 6: Rental market indicators

6 Conclusion

In this paper, we develop an agent-based model of the UK housing market. We demonstrate the usefulness of this approach for research and policy by exploiting the heterogeneous nature of households along multiple dimensions, including their income, their home ownership status or their characteristics as borrowers, among others. This allows us to identify which categories of agents may be particularly affected by different housing policy experiments. This level of granularity is informative as it facilitates the identification of unintended side effects of policies, the analysis of spillover effects to other markets segments, and the assessment of the full

distributions of variables of interest.

While the model allows for a rich analysis of the housing market and the measures affecting it, we limit ourselves, in the interest of space, to two policy experiments, a hard loan-to-value (LTV) limit and a soft loan-to-income (LTI) limit. We highlight a range of effects these two policy experiments have on the housing market, with a particular focus on house price booms and busts as well as sub-groups of households such as first-time buyers (FTBs), home movers (HMs), and renters.

Among the several results presented in this paper, three findings stand out. First, the housing policy experiments tend to mitigate the house price cycle. That is, variables such as housing transactions, mortgage approvals, and house prices tend to be lower during booms but higher during busts compared with the benchmark of no policy. The reason is that the availability of credit and therefore leverage is reduced, leaving the market less volatile. Second, a given policy can impact several different risk metrics, including the ones it does not target directly. For example, an LTI limit can also reduce the LTV ratios available for the same group of borrowers. This effect should be taken into account when calibrating any housing market policy to achieve the desired overall reduction in risk. Third, the policy experiments affect different types of households differently. Given the nature of the policy experiments performed in this paper (i.e. targeting owner-occupiers but not the BTL sector), BTL investors tend to benefit at the expense of owner-occupiers. FTBs are generally more negatively impacted than HMs, for example by losing some access to high-risk mortgages. These compositional changes also spill over to the rental market. This happens through two channels: the policies prevent some FTBs and HMs from securing sufficiently large mortgages to purchase properties, therefore increasing demand in the rental sector. The properties are instead absorbed by the BTL sector, which increases the supply of houses on offer in the rental market. The policy experiments therefore lead to a compositional shift in home ownership away from the owner-occupier to the BTL sector. Based on our calibration for the UK, demand for rental properties increases by more than their supply in case of the LTV policy experiment, thereby raising both the number of rental transactions and average rents. In case of the LTI policy experiment, the increase in supply and demand is more balanced, leading to a smaller increase in rental transactions and a very small decrease in

rents.

We have left several potential extensions of the model to future research. First, the housing model should be embedded in a macro framework in order to be able to assess the model dynamics in a general equilibrium setting. This could be achieved by merging the model into one of the ABM families described in the literature review above. In particular, the lack of monetary policy in our model limits the role of the interest rate as an endogenous determinant of house price and credit cycles (even though there is an ad hoc countercyclical interest rate mechanism present). Existing agent-based models (ABMs) that assess the role of monetary policy in a macro framework include, for instance, [Popoyan et al. \(2017\)](#) and [Alexandre & Lima \(2020\)](#). Furthermore, incorporating a more detailed banking sector into the model could shed light on understanding heterogeneous responses of banks, their risk-taking incentives and balance-sheet adjustments in response to macroprudential housing policies. This would complement, for example, the empirical analysis by [Acharya et al. \(2020\)](#) who find that the LTV and LTI limits introduced in Ireland in 2015 resulted in lenders adjusting their portfolios and increasing their risk exposure to corporate credit and holdings of securities. Similarly, [Ely et al. \(2021\)](#) present evidence that macroprudential policies, including LTV and debt-to-income ratios, can mitigate bank risk as measured by the Z-score, mainly by limiting leverage. Furthermore, [Peydró et al. \(2020\)](#) show that as a response to the implementation of an LTI flow limit in the UK in 2014, constrained lenders decreased the proportion of their high-LTI lending, and this reduction affected the low income borrowers more.

Second, we have only considered two types of policy experiments in this paper. Future research could focus on other measures such as debt-service-to-income limits, tax policies, or policies conditional on the state of the housing market. A growing literature has focused on the impact of macroprudential bank capital requirements, such as [Cincotti et al. \(2012b\)](#), [Krug et al. \(2015\)](#), [Popoyan et al. \(2017\)](#), [Van Der Hoog & Dawid \(2019\)](#), [Alexandre & Lima \(2020\)](#), and [Popoyan et al. \(2020\)](#), though to our knowledge no studies have comprehensively assessed a combination of monetary policy as well as macroprudential capital and housing instruments. Other policies could comprise credit quotas and credit rationing, as discussed by [van der Hoog \(2018\)](#).

Third, we have not explicitly considered how to calibrate policies optimally in our analysis. A starting point could be to specify different objective functions of a policymaker based on observable variables and determine the best combination of policies which maximise the value of the respective objective function. A welfare criterion is relatively difficult to define in ABMs, but could provide a worthwhile subject of future research.

A final potential avenue of research would be to consider spatial features explicitly in the model. In particular, expanding it to include different regions, with their specific house price dynamics, would allow one to study the heterogeneous impact of a given macroprudential policy across these regions. Furthermore, this improvement would open the way for an exhaustive analysis of affordability, taking into account spatial heterogeneities in the distributions of income, wealth and housing prices.

Acknowledgements and disclaimer

As mentioned in the Introduction, this paper grew out of a Bank of England Staff Working Paper (#619), and we would particularly like to thank our collaborators Rafa Baptista, Katie Low, and Daniel Tang. We would also like to thank Robert Axtell, Marco Bardoscia, Sebastian Barnes, Philippe Bracke, Leif Brubakk, Chris Carroll, Gennaro Catapano, Graeme Cokayne, Allan Davids, Co-Pierre Georg, Andy Haldane, William Hynes, Sujit Kapadia, Junior Maih, Fergal McCann, Bence M  r  , John Muellbauer, Mauro Napoletano, Marco Pangallo, Donovan Platt, Lilit Popoyan, Mikhail Prokopenko, David Pugh, Andrea Roventini, Ruben Tarne, Nikolett V  g  , Matthew Waldron, and Gavin Wallis for their support and valuable comments. Georgia Papamarkidou provided excellent research assistance. We would also like to express our gratitude to the participants of WEHIA 2017 and the 2017 Bank of England and Office of Financial Research joint conference as well as workshops at the Bank of England, the Central Bank of Hungary, the Central Bank of Ireland, Bank of Russia, Norges Bank, OECD, Essex University, and the Alan Turing Institute for their valuable comments. Finally, Adrian Carro was partially supported by the UK Engineering and Physical Science Research Council under grant EP/N017064/1: Multi-scale Infrastructure Systems Analytics (MISTRAL).

Disclaimer: Any views expressed in this article are solely those of the authors and should

not be attributed to the Banco de España, the Eurosystem or the Bank of England or their policy committees. The analysis presented in this article uses HM Land Registry data © Crown copyright and database right 2021. This data is licensed under the Open Government Licence v3.0.

References

- Acharya, V. V., Bergant, K., Crosignani, M., Eisert, T., & McCann, F. J. (2020). *The anatomy of the transmission of macroprudential policies*. Working Paper 27292, NBER.
- Adelino, M., Schoar, A., & Severino, F. (2012). *Credit supply and house prices: evidence from mortgage market segmentation*. Working Paper 17832, NBER.
- Alexandre, M. & Lima, G. T. (2020). Combining monetary policy and prudential regulation: an agent-based modeling approach. *Journal of Economic Interaction and Coordination*, 15(2), 385–411.
- Aron, J. & Muellbauer, J. (2016). Modelling and forecasting mortgage delinquency and foreclosure in the UK. *Journal of Urban Economics*, 94, 32–53.
- Ashraf, Q., Gershman, B., & Howitt, P. (2017). Banks, market organization, and macroeconomic performance: an agent-based computational analysis. *Journal of Economic Behavior & Organization*, 135, 143–180.
- Assenza, T., Gatti, D. D., & Grazzini, J. (2015). Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control*, 50, 5–28.
- Association of Residential Letting Agents (2013). *ARLA Members Survey of the Private Rented Sector (Fourth Quarter 2013)*. Technical report, ARLA.
- Axtell, R., Farmer, D., Geanakoplos, J., Howitt, P., Carrella, E., Conlee, B., Goldstein, J., Hendrey, M., Kalikman, P., Masad, D., Palmer, N., & Yang, C.-Y. (2014). An agent-based model of the housing market bubble in metropolitan washington, DC. In *Deutsche Bundesbank's*

Spring Conference on “Housing markets and the macroeconomy: Challenges for monetary policy and financial stability (pp. 5–6).

Bank of England (2018). *A loan-level data collection for buy-to-let lending, 2017-2018*. [data collection], Bank of England (BoE).

Bank of England (2019). *Financial Stability Report - December 2019*. Financial Stability Report, Bank of England.

Bank of England/NMG (2018). *Bank of England/NMG household survey data, 2014-2018*. [data collection], Bank of England (BoE).

Baptista, R., Farmer, J. D., Hinterschweiger, M., Low, K., Tang, D., & Uluc, A. (2016). *Macroprudential policy in an agent-based model of the UK housing market*. Staff Working Paper 619, Bank of England.

Caiani, A., Godin, A., Caverzasi, E., Gallegati, M., Kinsella, S., & Stiglitz, J. E. (2016). Agent based-stock flow consistent macroeconomics: Towards a benchmark model. *Journal of Economic Dynamics and Control*, 69, 375–408.

Catapano, G., Franceschi, F., Loberto, M., & Michelangeli, V. (2021). *Macroprudential policy analysis via an agent based model of the real estate sector*. Temi di Discussione (Working Paper) 1338, Bank of Italy.

Chen, Z. & Lux, T. (2018). Estimation of sentiment effects in financial markets: A simulated method of moments approach. *Computational Economics*, 52(3), 711–744.

Cincotti, S., Raberto, M., & Teglio, A. (2012a). The eurace macroeconomic model and simulator. In M. Aoki, K. Binmore, S. Deakin, & H. Gintis (Eds.), *Complexity and Institutions: Markets, Norms and Corporations* (pp. 81–106). Palgrave Macmillan UK.

Cincotti, S., Raberto, M., & Teglio, A. (2012b). Macroprudential policies in an agent-based artificial economy. *Revue de l’OFCE*, 2012/5 (No 124), 205–234.

- Cokayne, G. (2019). *The Effects of Macro-prudential Policies on House Price Cycles in an Agent-based Model of the Danish Housing Market*. Working Paper 138, Danmarks Nationalbank.
- Dawid, H. & Delli Gatti, D. (2018). Agent-Based Macroeconomics. *Handbook of Computational Economics*, 4, 63–156.
- Dawid, H., Gemkow, S., Harting, P., Van der Hoog, S., & Neugart, M. (2012). The Eurace@Unibi Model: An Agent-Based Macroeconomic Model for Economic Policy Analysis. *Bielefeld Working Papers in Economics and Management*, No. 05-2012.
- Dawid, H., Harting, P., & Neugart, M. (2014). Economic convergence: Policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control*, 44, 54–80.
- Dawid, H., Harting, P., Van der Hoog, S., & Neugart, M. (2019). Macroeconomics with heterogeneous agent models: fostering transparency, reproducibility and replication. *Journal of Evolutionary Economics*, 29(1), 467–538.
- Delli Gatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., & Gallegati, M. (2011). *Macroeconomics from the Bottom-up*, volume 1. Springer Science & Business Media.
- Department for Communities and Local Government (2013). *English Housing Survey, 2011-2012: Household Data*. [data collection]. SN: 7362, UK Data Service.
- Department for Environment, Food and Rural Affairs, Office for National Statistics (2020). *Living Costs and Food Survey, 2011 (2nd Edition)*. [data collection]. SN: 7272, UK Data Service.
- Dosi, G., Fagiolo, G., & Roventini, A. (2010). Schumpeter meeting keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control*, 34(9), 1748–1767.
- Dosi, G. & Roventini, A. (2019). More is different... and complex! the case for agent-based macroeconomics. *Journal of Evolutionary Economics*, 29(1), 1–37.

- Ely, R. A., Tabak, B. M., & Teixeira, A. M. (2021). The transmission mechanisms of macro-prudential policies on bank risk. *Economic Modelling*, 94, 598–630.
- Evans, B. P., Glavatskiy, K., Harré, M. S., & Prokopenko, M. (2021). The impact of social influence in Australian real estate: market forecasting with a spatial agent-based model. *Journal of Economic Interaction and Coordination*, (pp. 1–53).
- Fabretti, A. (2013). On the problem of calibrating an agent based model for financial markets. *Journal of Economic Interaction and Coordination*, 8(2), 277–293.
- Fagiolo, G., Guerini, M., Lamperti, F., Moneta, A., & Roventini, A. (2019). Validation of agent-based models in economics and finance. In *Computer Simulation Validation* (pp. 763–787). Springer.
- Fagiolo, G. & Roventini, A. (2017). Macroeconomic Policy in DSGE and Agent-Based Models Redux: New Developments and Challenges Ahead. *Journal of Artificial Societies and Social Simulation*, 20(1).
- Farmer, J. D. & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256), 685–686.
- Favara, G. & Imbs, J. (2015). Credit supply and the price of housing. *American Economic Review*, 105(3), 958–92.
- Filatova, T. (2015). Empirical agent-based land market: Integrating adaptive economic behavior in urban land-use models. *Computers, environment and urban systems*, 54, 397–413.
- Financial Conduct Authority (2018). *Mortgage Product Sales Data, 2005-2018*. [data collection], Financial Conduct Authority (FCA).
- Franke, R. (2009). Applying the method of simulated moments to estimate a small agent-based asset pricing model. *Journal of Empirical Finance*, 16(5), 804–815.
- Franke, R. & Westerhoff, F. (2012). Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 36(8), 1193–1211.

- Ge, J. (2013). Who creates housing bubbles? an agent-based study. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation* (pp. 143–150).: Springer.
- Ge, J. (2017). Endogenous rise and collapse of housing price: an agent-based model of the housing market. *Computers, Environment and Urban Systems*, 62, 182–198.
- Geanakoplos, J., Axtell, R., Farmer, J. D., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N. M., & Yang, C.-Y. (2012). Getting at Systemic Risk via an Agent-Based Model of the Housing Market. *American Economic Review*, 102(3), 53–58.
- Gilli, M. & Winker, P. (2003). A global optimization heuristic for estimating agent based models. *Computational Statistics & Data Analysis*, 42(3), 299–312.
- Gillingham, R. (1980). Estimating the user cost of owner-occupied housing. *Monthly Labor Review*, 103(2), 31–35.
- Glavatskiy, K. S., Prokopenko, M., Carro, A., Ormerod, P., & Harre, M. (2021). Explaining herding and volatility in the cyclical price dynamics of urban housing markets using a large-scale agent-based model. *SN Business & Economics*, 1(6), 1–21.
- Guerrero, O. A. (2020). Decentralized markets and the emergence of housing wealth inequality. *Computers, Environment and Urban Systems*, 84, 101541.
- Haas, A. & Jaeger, C. (2005). Agents, bayes, and climatic risks-a modular modelling approach. *Advances in Geosciences*, 4, 3–7.
- Haldane, A. G. & Turrell, A. E. (2019). Drawing on different disciplines: macroeconomic agent-based models. *Journal of Evolutionary Economics*, 29(1), 39–66.
- Huang, Q., Parker, D. C., Sun, S., & Filatova, T. (2013). Effects of agent heterogeneity in the presence of a land-market: A systematic test in an agent-based laboratory. *Computers, environment and urban systems*, 41, 188–203.
- Jordà, Ò., Schularick, M., & Taylor, A. M. (2016). The great mortgaging: housing finance, crises and business cycles. *Economic Policy*, 31(85), 107–152.

- Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3), 697–743.
- Krug, S., Lengnick, M., & Wohltmann, H.-W. (2015). The impact of basel iii on financial (in) stability: an agent-based credit network approach. *Quantitative Finance*, 15(12), 1917–1932.
- Laliotis, D., Buesa, A., Leber, M., & Población, J. (2020). An agent-based model for the assessment of LTV caps. *Quantitative Finance*, 20(10), 1721–1748.
- Mandel, A., Jaeger, C., Fuerst, S., Lass, W., Lincke, D., Meissner, F., Pablo-Marti, F., & Wolf, S. (2010). *Agent-based dynamics in disaggregated growth models*. Working Paper 2010.77, CES.
- Mérő, B. & Vágó, N. (2018). Keresletvezérelt lakáspiaci modell a lakáshitelezést szabályozó makroprudenciális eszközök tanulmányozására. *Közgazdasági Szemle*, 65(11), 1115–1153.
- Ministry of Housing, Communities and Local Government (2019). *English Private Landlord Survey, 2018*. [data collection]. SN: 8455, UK Data Service.
- Moneyfacts Group PLC (2011). *MoneyFacts Analyser: Mortgage Products, 2011*. [data collection], Moneyfacts Group PLC.
- Office for National Statistics (2009). *Wealth in Great Britain, main results from the Wealth and Assets Survey 2006/2008*. Technical report, Colchester, Essex: UK Data Archive.
- Office for National Statistics, Social Survey Division (2022). *Wealth and Assets Survey, Waves 1-5 and Rounds 5-7, 2006-2020 (15th Edition)*. [data collection]. SN: 7215, UK Data Service.
- Ozel, B., Nathanael, R. C., Raberto, M., Teglio, A., & Cincotti, S. (2019). Macroeconomic implications of mortgage loan requirements: an agent-based approach. *Journal of Economic Interaction and Coordination*, 14(1), 7–46.
- Pangallo, M., Nadal, J.-P., & Vignes, A. (2019). Residential income segregation: A behavioral model of the housing market. *Journal of Economic Behavior & Organization*, 159, 15–35.

- Parker, D. C. & Filatova, T. (2008). A conceptual design for a bilateral agent-based land market with heterogeneous economic agents. *Computers, Environment and Urban Systems*, 32(6), 454–463.
- Peydró, J.-L., Rodriguez-Tous, F., Tripathy, J., & Uluc, A. (2020). *Macroprudential policy, mortgage cycles and distributional effects: Evidence from the UK*. Bank of England working papers 866, Bank of England.
- Platt, D. & Gebbie, T. (2018). Can agent-based models probe market microstructure? *Physica A: Statistical Mechanics and its Applications*, 503, 1092–1106.
- Popoyan, L., Napoletano, M., & Roventini, A. (2017). Taming macroeconomic instability: Monetary and macro-prudential policy interactions in an agent-based model. *Journal of Economic Behavior & Organization*, 134, 117–140.
- Popoyan, L., Napoletano, M., & Roventini, A. (2020). Winter is possibly not coming: Mitigating financial instability in an agent-based model with interbank market. *Journal of Economic Dynamics and Control*, 117, 103937.
- Seppecher, P., Salle, I. L., & Lang, D. (2019). Is the market really a good teacher? *Journal of Evolutionary Economics*, 29(6), 299–335.
- Seppecher, P., Salle, I. L., & Lavoie, M. (2018). What drives markups? Evolutionary pricing in an agent-based stock-flow consistent macroeconomic model. *Industrial and Corporate Change*, 27(6), 1045–1067.
- Tarne, R., Bezemer, D., & Theobald, T. (2021). *The Effect of Borrower-Specific Loan-to-Value Policies on Household Debt, Wealth Inequality and Consumption Volatility: An Agent-Based Analysis*. IMK Working Paper 212, Hans-Böckler-Stiftung.
- van der Hoog, S. (2018). The limits to credit growth: Mitigation policies and macroprudential regulations to foster macrofinancial stability and sustainable debt. *Computational Economics*, 52(3), 873–920.

- Van Der Hoog, S. & Dawid, H. (2019). Bubbles, crashes, and the financial cycle: The impact of banking regulation on deep recessions. *Macroeconomic Dynamics*, 23(3), 1205–1246.
- WhenFresh (2016). *Zoopla data, 2003-2015*. [data collection], WhenFresh/Zoopla.
- Wolf, S., Fürst, S., Mandel, A., Lass, W., Lincke, D., Pablo-Marti, F., & Jaeger, C. (2013). A multi-agent model of several economic regions. *Environmental modelling & software*, 44, 25–43.
- Wong, E., Fong, T., fai Li, K., & Choi, H. (2011). *Loan-to-Value Ratio as a Macro-Prudential Tool - Hong Kong's Experience and Cross-Country Evidence*. Working Papers 1101, Hong Kong Monetary Authority.
- Yun, T.-S. & Moon, I.-C. (2020). Housing Market Agent-Based Simulation with Loan-To-Value and Debt-To-Income. *Journal of Artificial Societies and Social Simulation*, 23(4), 5.

Appendices

Appendix A Detailed Model Description

As in the main text, we use here Latin letters for identifying specific model variables and Greek letters for parameters, with the exception of ε , which denotes a random draw from a normal distribution with mean ε_μ and standard deviation ε_σ , η , which denotes a random draw from a distribution estimated from data, and σ , which denotes the standard logistic function $\sigma(x) = 1/(1 + e^{-x})$. The specific values of the parameters, as well as a description of the estimation or calibration procedure and the data sources used, can be found in Appendix B by referring to the corresponding subsection of Section 3 or equation number.

A.1 Initial model set-up

At the beginning of each simulation, the target number of households are created, drawing all relevant household characteristics from suitable distributions. Notably, households are assigned an initial age, an income percentile, a propensity to save, an initial financial wealth and, potentially, a buy-to-let (BTL) flag and an investor type. Households are initially set to live in social housing, as no house has yet been created. See Appendix A.2 for further details on the initialisation of the different household variables when a household is created, as well as on the concept of social housing.

In a similar manner, a number of houses are also created at the beginning of the simulation, corresponding to the target ratio of houses to households and remaining constant thereafter. This stock of houses is distributed at random among existing households, who will move into the first house they are given. Owner-occupier households —without the BTL flag— receiving more than one house will put the resulting extra properties for sale while also temporarily offering them for rent. Investor household —with the BTL flag— receiving more than one house will decide, for each of them, whether to keep it and add it to their portfolio or sell it, in which case they will still temporarily offer it for rent. As a consequence of this initialisation assumption —random and free initial distribution of houses—, a “spin-up” period is needed at the beginning

of each simulation, such that a realistic assignment of houses and mortgages to households can be reached. While the model output typically stabilises around realistic values in about 200 months, we allow the model to spin up for 500 months. For the purposes of processing the simulation output, the spin-up part of the simulation is always completely discarded.

Finally, a number of model variables requiring a pre-simulation value, such as exponential moving averages, are initialised as estimated from data (see Appendix B). In particular, initial values are set for the exponential moving average price for each quality segment and for each of the two markets. These are estimated from the corresponding real price distributions, assuming each quality segment simply maps onto a corresponding percentile of the real price distribution. In this way, even if houses are initially allocated at random among the households and given for free (as if inherited), both initial sale and rental prices—as well as their distributions—will be close to the real ones, given the backward-looking ask price setting mechanism of the model (see Section 3.3).

A.1.1 House qualities

Each house is characterised by a single permanent parameter, its quality, which is assigned uniformly at random when the house is created from a discrete set of possible integer values, with higher values implying higher quality. The purpose of this quality parameter is to serve as a proxy for all possible features making houses different and, importantly, some more desirable than others, such as their location, size, condition or dwelling type. Due to its uniform random assignment and its discrete character, this parameter will effectively divide the market into approximately equally sized market segments, each corresponding to a given quality from its set of possible values. Within each of these segments, houses are indistinguishable but for their price. Across segments, on the contrary, houses are not only distinguishable but, most importantly, they can be ordered according to their quality values. In fact, this reduction of all house features into a single dimension allows for all households to have the exact same order of preferences regarding a given set of houses. This significantly simplifies the internal functioning of the market mechanism (see Subsection 3.5), as a single ordered list of preferences is needed instead of a different one for each bidder. Moreover, it allows prices to differentiate between the

various segments, with higher quality segments being characterised by higher prices than lower quality segments. As a consequence, the former are accessible only to wealthier households while the latter receive bids mostly by less affluent households, given the general preference for higher qualities. Note, however, that this is a non-negligible simplification from the intrinsic complexity of real housing markets, where different households might prefer different locations, different types of dwellings or even different combinations of house features. Thus, our approximation would work better the more amenable the features under consideration are to an objective and universal order of preferences. For example, at a given price level, it is usual to prefer a larger rather than a smaller house, or a better built rather than a worse built one.

Bearing in mind the market mechanism of the model, described in Subsection 3.5, it is easy to see that, when the number of different quality segments considered is very small, there will be many failed bid-offer matches and thus fewer sales. This is due to the fact that, with fewer quality segments, it is harder to correctly order bidders and sellers, and so some wealthy households will end up buying cheap houses which are then not available for poorer households. In other words, fewer quality segments mean more competition for the same resources. On the contrary, when the number of quality segments is very large, there is very rarely any sale per quality segment per time step, and so market statistics suffer strong fluctuations. In view of this, we choose to have the maximum number of quality segments compatible with having at least, on average, one sale per quality segment per time step. To this end, we approximate the number of sales per month as the monthly sale probability for owner-occupiers multiplied by the total number of houses, which is itself equal to the number of households multiplied by the number of houses per household. Note that this assumes that the selling probability of owner-occupiers is a good approximation for assessing the number of new sellers per month, which is coherent with the fact that there are many more owner-occupiers than BTL investors. Furthermore, the approximation assumes that the market does not grow over time, such that the number of new sellers is similar to the number of sales. Given the fixed population size of the model, this must, on average, be the case. Finally, note that this choice for the number of quality segments ensures a proper scaling with the number of households in the model.

A.2 Demographics

The main goal of the demographic component of the model is to keep a constant distribution of age and a total number of households fluctuating around a constant target while households age. In other words, we are not trying to reproduce a realistic demographic path, with its associated population growth and progressive inversion of the age pyramid. On the contrary, we are trying to keep as many demographic characteristic as possible fixed at their values corresponding to a certain point in history (2011) while allowing for a dynamic population with household renewal and ageing. In order to accomplish this with the least number of births and deaths, we simply choose a birth rate dependent on age so as to counteract the declining population in lower age bins as households age and a death rate dependent on age so as to counteract the growing population in higher age bins as households age. The specific reasons that can lead to the creation of a new household, from children leaving their parents' home to splitting due to a divorce, as well as those leading to the dissolution of an existing one, from merging due to marriage to the actual death of all the household members, are all beyond the scope of this contribution.

New households are born with all their characteristics drawn from suitable distributions. In particular, households are assigned:

1. An initial age of the Household Reference Person (HRP). This is drawn uniformly at random within the edges of the age bin for which the household needs to be created, and it will then naturally evolve in the course of the simulation, driving the evolution of other household attributes. The concept of HRP, used in a number of government surveys, refers to the person within the household whose personal characteristics are used for characterising the household. In particular, the HRP is defined in the Wealth and Assets Survey ([Office for National Statistics, 2009](#)) as follows: in households with a sole householder, that person is the HRP, while in households with joint householders, the person with the highest income is taken as the HRP. If more than one householder have exactly the same income, then the oldest one is chosen as the HRP.
2. A permanent income percentile. This is assigned uniformly at random and it will remain

fixed throughout the household’s lifespan. Together with the household’s age, it will determine the household’s income, which will then evolve in time driven by age (see Section 3.2 for further details).

3. A permanent propensity to save. This is assigned uniformly at random and it will remain fixed throughout the household’s lifespan. Together with the household’s income, it will determine the household’s target financial wealth. This target will then evolve in time driven by income —and thus driven by age— and it will be used for consumption and saving decisions (see Appendix A.3 for further details).
4. An initial financial wealth. This is simply chosen as the target financial wealth corresponding to the household’s permanent propensity to save and its initial income, itself derived from the household’s permanent income percentile and its initial age (see Appendix A.3 for further details).

Apart from these features, which are set for all households, a permanent buy-to-let (BTL) flag is randomly added to some of them, with a probability dependent on income percentile. This flag signals the household’s willingness to invest in BTL properties. That is, only households with this flag can ever buy additional properties beyond their homes to offer them on the rental market. Note, however, that, due to inheritances (see below), households without this flag can temporarily hold and rent out additional properties while they attempt to sell them. Regarding the probability of a household being assigned the BTL flag, since we only have data on the fraction of households per income percentile receiving any rental income (i.e., on the fraction of households actually renting properties out), we introduce a multiplier or adjustment parameter in order to account for households willing to do it but not yet having managed to acquire and rent out any investment property (see Appendix B for further calibration details). Finally, households receiving the BTL flag are also assigned an investor type, to be randomly chosen —with probabilities according to survey data— among capital-gains-driven, rental-income-driven or mixed type. The investor type is associated with a specific intensity of the household’s interest in capital gains as opposed to rental yield.

All new households are born into social housing, which we define as a temporary and free

accommodation while they try to find a house to rent or buy. In general, this concept of social housing can be thought of as representing a variety of accommodation options going beyond the scope of the model, which focuses only on the privately sold and privately rented sectors. Particular examples of such situations are homelessness, living with parents or living out of housing and other social benefits. Importantly, households never choose to be in social housing, but are put there if they fail to secure any other form of housing or as a temporary state between selling a house and buying or renting a new one. Since no housing payment is deducted from the households' income while in social housing, one can think of it as a very simple form of in-kind housing benefit.

When a household dies, another randomly chosen household inherits all its properties and positive financial wealth net of mortgage debt. If the deceased household had any houses on any market, these are first taken off the market. If it had any investment properties rented out, then any tenants living in the houses are evicted. Outstanding mortgages are paid off as much as possible from the deceased household's financial wealth, with any further outstanding amount being written off. If the deceased household was renting, then the rental contract is simply terminated. Upon inheriting a house, *(i)* households currently in social housing will immediately move into it, *(ii)* renters will first terminate their rental contracts and then move into it, *(iii)* owner-occupiers will put the inherited house for sale while also temporarily offering it for rent and *(iv)* BTL investors will decide whether to add it to their portfolio or sell it, in which case they will still temporarily offer it for rent.

A.3 Desired consumption and financial wealth

Given the model's lack of macroeconomic dynamics and the strong influence of the financial wealth of households for their housing decisions, we model their desired consumption in a highly stylised manner, with the sole goal of leading to an accurate distribution of financial wealth among the population. In particular, we model a household's desired consumption C as

$$C = \min \left(\max \left(\frac{1}{2} (w - w_t + y_d), 0 \right), C_{\max} \right), \quad (16)$$

where y_d is the household's current disposable income (that is, after tax, essential non-housing consumption and housing expenses), w is its current financial wealth (after adding this month's disposable income), w_t is its target or desired level of financial wealth, as computed with the inverse cumulative distribution of financial wealth conditional on gross income analysed at the household's propensity to save, and C_{\max} is a maximum level of monthly desired consumption (see Appendix B Table 9 for estimation details). The maximum simply ensures that desired consumption is always positive or zero. Subtracting a monthly disposable income y_d from the target financial wealth w_t sets the former as the distance from which the household starts relaxing its saving behaviour, thus allowing for some desired, non-essential consumption even before the target is reached. This, together with the factor $1/2$, ensures that the household's financial wealth exponentially approaches its target level, both from below and from above. Furthermore, this functional form also ensures that households with higher incomes consume more.

A household's financial wealth can increase substantially when selling a house, bringing it far above its target level. In order to avoid unrealistically high levels of consumption in these cases, we set a maximum monthly desired consumption C_{\max} , which we define as a fraction of the household's gross annual income, estimated by focusing on the consumption levels of households in the 99th income percentile.

Finally, the resulting desired consumption is subtracted from the household's current financial wealth. Thus, households save if their disposable income is larger than their desired consumption and dissave otherwise, drawing down their financial wealth.

A.4 Bankruptcy

Since the model does not seek to capture the nuances of bankruptcy dynamics (from delinquency to foreclosure), we simply consider any household unable to meet a payment commitment as directly bankrupt. Whenever this happens, we make sure all committed payments are always made by artificially injecting as much cash as necessary, without further action against the bankrupt household. In any case, due to the absence of explicit income shocks,⁴⁴ bankrupt-

⁴⁴ Incomes can decrease in our model, but these decreases are mostly associated with retirement age, by which time all residential mortgages have been repaid.

cies are not frequent in the model.

Among all household expenses, only housing payments can lead to bankruptcies: taxes are a fraction of the household's income, essential non-housing consumption is estimated to be below the minimum net income, and non-essential non-housing consumption only affects the household's financial wealth in excess of its desired level, thus never leading to negative values. While in principle designed for households to be able to fulfil their payments, the housing decisions of a household can end up leading to payment commitments it cannot afford. In particular, there are three ways in which housing expenses can lead to bankruptcies:

1. Very low income renters might make rental bids too close to their maximum payment limit, which can then become unaffordable under the slightest decrease of their income. This combines with the fact that, as a result of the parameter values estimated from data, the evolution of the income of very low income percentile households is characterised by more fluctuations than higher percentiles. Thus, these household can experience some decreases of their income even before retirement age.
2. Buy-to-let investors unable to rent their properties out might be unable to meet their interest-only mortgage payments.
3. Buy-to-let investors without savings large enough to repay the principal on their properties and unable to sell them before maturity will default on this final principal repayment.

Note that these types of bankruptcy happen in the absence of any explicitly modelled income shock and, in particular, the latter two are a result of the model's market dynamics—respectively, of the rental and the sales market.

A.5 Bank interest rate spread

The interest rate paid by households on their mortgages is the sum of a fixed policy rate exogenously set by the Central Bank and a variable interest rate spread set by the bank. This spread compensates the bank for factors such as credit risk, administrative expenses, and a profit margin, but these factors are not modelled explicitly.

We model interest rate spread movements as a linear function of movements of the amount of credit per household. That is, we compute the mortgage interest rate spread for the next month based on the spread for this month and the resulting increase or decrease in the demand for credit per household, which in this model is equivalent to supply,

$$i_{t+1} = i_t + \varphi \frac{(C_t - C_{t-1})}{N}, \quad (17)$$

where i_t is the interest rate spread at time t , C_t is the total credit supply (or demand) at time t , N is the number of households and φ is a constant that we estimate from data as the average of the absolute value of month to month interest rate differences divided by the average of the absolute value of month to month amount of credit per household differences (see Appendix B Table 13 for estimation details). In this way, and bearing in mind that the policy rate is constant, increases (decreases) in the demand for credit will be counteracted with proportional increases (decreases) in interest rates, such that both the demand for credit and interest rates will tend to fluctuate around their initial values, which we set as estimated from data.

A.6 Market bid-up process

When a given house offered in any of the two housing markets —sales or rental— is matched with more than one bidder, the seller or property owner bids up the price k times by multiplying it k times by a bid-up factor larger than 1 and estimated from data (see Appendix B Table 14 for estimation details),

$$p_s^{\text{new}} = \chi^k p_s, \quad (18)$$

p_s^{new} is the updated offer price after the bid-up, where p_s is the old offer price and χ is the bid-up factor. The idea behind bidding up the price k times is to make this k dependent on the number of bidders, such that properties with many bidders get to increase their price more than properties with just a few bidders. We model this dependency by assuming that the bids arrive at random days during the month and that a bid-up will take place whenever a new bid arrives within three days of the previous one. Thus, the bid-up process stops whenever a bid remains as the latest one for more than three days. Mathematically, this is equivalent to drawing, for each

offer with multiple bids, a value of k from the geometric distribution

$$P(k) = (1 - P_{\text{success}})^k P_{\text{success}} , \quad (19)$$

with

$$P_{\text{success}} = \left(1 - \frac{3}{30}\right)^{b-1} , \quad (20)$$

where b is the logarithm (base 10) of the number of bids received for this offer.⁴⁵ Note that k has a certain probability of being $k = 0$, thus leading to no bid-up.

A.7 Market price information

Households make use of market price information when making their housing decisions. Two types of market information are available for the sales and rental markets, respectively: the house price index (HPI) and rental price index (RPI) accounting for the behaviour of the whole market, and the exponential moving average sale and rental price for each market quality segment. In order to build a market index for each market, we follow a mix adjustment methodology. In particular, for a given month, the market index is defined as the average price over all transactions completed during that month divided by the average reference price over the same set of transactions. The reference price of a house is the average price of houses of that quality as estimated from data, and can be thought of as a measure of that quality. The main reason for using exponential moving average sale and rental prices instead of simple averages is that, due to the scaling down of the population, the number of transactions in a given month may be quite small for some of the market quality segments.⁴⁶

When trying to assess the mark-to-market value of a house, it is reasonable to think that households consider both the general market behaviour, as captured by the market indices described above, and the behaviour of the specific market quality segment under consideration,

⁴⁵ The design of the market mechanism, which leads to all households able to afford a certain quality bidding for the cheapest house within that quality, frequently leads to rather large numbers of bids matched to the same offer. Thus, in order to get realistic numbers of bidders for a given property, we take the logarithm (base 10) of the number of matched bids for computing bid-ups.

⁴⁶ Following one of the standards in the literature, we use a decay constant corresponding to assigning a cumulative weight of $1/4$ to events older than a year and thus to a half-life of half a year.

that is, the subset of houses of the same quality as the one being assessed. In order to capture this, we introduce the following modification of the exponential moving average sale and rental price, respectively, for each market quality segment used by households to make their decisions,

$$p'_Q = \rho p_Q + (1 - \rho) I p_Q^{\text{ref}}, \quad (21)$$

where p_Q is the exponential moving average sale or rental price for houses of quality Q , I is the house or rental price index for this month, p_Q^{ref} is the reference price of houses of quality Q and ρ is a parameter setting the relative weight given to the behaviour of the market quality segment in particular as opposed to the general market behaviour.

Appendix B Estimation and calibration of model parameters

The parameters of the model can be classified into three different categories depending on the type of estimation or calibration procedure used to set their value:

1. **Estimated parameters:** Those parameters and distributions which can be directly estimated from available data sources. Most parameters in our model (56) belong to this category. We use data from 2011, where available, for reasons set out in Section 4. If data from 2011 are not available, we choose the closest available year.
2. **Postulated parameters with a sensitivity analysis:** Those which cannot be directly estimated from readily available data sources but for which plausible values can be postulated. Furthermore, their effects are contained to specific parts of the model, which makes them amenable to an independent sensitivity analysis, rather than a full calibration. This applies to 6 of the parameters in our model.
3. **Calibrated parameters:** Those which cannot be directly estimated from readily available data sources, whose effects are widespread across the model and for which it is unfeasible to postulate plausible values. This applies to 5 parameters in our model and, in particular, we use the method of simulated moments for their calibration.

In the following tables we report the parameter values and how they are derived. For estimated parameters as described in point 1 above, we provide the underlying data sources. The list of data sources and their acronyms can be found in Table 7. Postulated parameters are labeled as “postulated”. Calibrated parameters are referred to as “calibrated”, and we provide more details on the calibration method in Appendix B.2. A complete description of all parameters is available in the Online Appendix.

B.1 Parameters

Each of the following tables corresponds to the respective subsections of Section 3.

Acronym	Acronym Description
ARLA	Association of Residential Letting Agents
ARLA Survey	Members Survey of the Private Rented Sector (collected by ARLA, 2013)
BoE	Bank of England
BoE-BTL	Bank of England buy-to-let data (collected by BoE, 2018)
DCLG	Department for Communities and Local Government
DEFRA	Department for Environment, Food and Rural Affairs
DWP	Department for Work and Pensions
EHS	English Housing Survey (collected by DCLG, 2013)
EPLS	English Private Landlord Survey (collected by MHCLG, 2019)
FCA	Financial Conduct Authority
HMRC	HM Revenue & Customs
Land Registry	HM Land Registry Price Paid Data
LCFS	Living Costs and Food Survey (collected by DEFRA/ONS, 2020)
MHCLG	Ministry of Housing, Communities and Local Government
MoneyFacts	Mortgage Products Analysers (collected by MoneyFacts, 2011)
NMG	BoE/NMG Household Survey (collected by BoE/NMG, 2018)
ONS	Office for National Statistics
PSD	FCA Mortgage Product Sales Data (collected by FCA, 2018)
WAS	Wealth and Assets Survey (collected by ONS, 2022)
Zoopla	WhenFresh/Zoopla Data (collected by WhenFresh, 2016)

Table 7: List of data sources

Parameter	Equation	Value	Source
Number of UK dwellings	-	22,626,000	ONS
Number of households in the UK	(17)	$N = 26,442,100$	ONS

Table 8: Initial model set-up

Parameter	Equation	Value	Source
Age	-	Estimated distribution	WAS
Gross monthly income	-	Estimated distribution	WAS
Financial wealth	(16)	Estimated distribution	WAS
Income tax	-	Rates by income bands	HMRC
National Insurance contributions	-	Rates by income bands	HMRC
Minimum government income support	-	£445.80	DWP
Essential non-housing consumption fraction	-	0.66	LCFS
Maximum desired consumption fraction	(16)	0.17	LCFS

Table 9: Income, non-housing consumption and financial wealth

Parameter	Equation	Value	Source
Desired purchase price			
Scale	(1)	$\alpha = 42.9036$	PSD, Land Registry
Exponent	(1)	$\beta = 0.7892$	PSD, Land Registry
Mean of normal noise	(1)	$\varepsilon_\mu = -0.0177$	PSD, Land Registry
Std. dev. of normal noise	(1)	$\varepsilon_\sigma = 0.4104$	PSD, Land Registry
Distribution of house prices			
Scale	(21)	12.1186	Land Registry
Exponent	(21)	0.6414	Land Registry
Distribution of rental prices			
Scale	(21)	6.2647	NMG
Exponent	(21)	0.6353	NMG
Weight on market vs segment house prices	(21)	$\rho = 0.5$	Calibrated
House price growth expectation			
Multiplier factor	(3)	$\gamma = 0.44$	NMG, Land Registry
Constant	(3)	$\zeta = -0.007$	NMG, Land Registry
Time period (years)		2	Postulated
Psychological cost of renting	(4)	$\lambda = 0.4$	Calibrated
Sensitivity parameter	(4)	$\theta = 0.001$	Calibrated
Desired rental price			
Scale	(5)	$\mu = 17.2166$	NMG
Exponent	(5)	$\nu = 0.3464$	NMG
Desired down-payment			
First-time buyers, Scale	-	10.35	PSD
First-time buyers, Shape	-	0.898	PSD
Home movers, Scale	-	11.15	PSD
Home movers, Shape	-	0.958	PSD

Table 10: Housing decisions if in social housing

Parameter	Equation	Value	Source
Holding period of owner-occupied houses	-	17 years	EHS
Initial sale price mark-up, η	(6)	Estimated distribution	Zoopla
Sale price reduction			
Monthly probability	-	0.0703	Zoopla
Mean percentage reduction	-	1.4531	Zoopla
Std. dev. of percentage reduction	-	0.7070	Zoopla

Table 11: Housing decisions as an owner-occupier

Parameter	Equation	Value	Source
Probability to receive BTL flag			
Raw probability	-	Estimated dist.	WAS
Probability adjustment	-	1.76	Calibrated
Buy-to-let motivations			
Rental-income-driven, probability	-	0.4927	NMG
Rental-income-driven (RI), weight	(7) and (10)	$\delta_{RI} = 0.1$	Postulated
Capital-gains-driven, probability	-	0.1458	NMG
Capital-gains-driven (CG), weight	(7) and (10)	$\delta_{CG} = 0.9$	Postulated
Mixed, probability	-	0.3615	NMG
Mixed (M), weight	(7) and (10)	$\delta_M = 0.5$	Postulated
Sensitivity of buy and sell decisions	(8) and (11)	$\xi = 100$	Calibrated
Down-payment fraction			
Mean	-	0.34	BoE-BTL
Standard deviation	-	0.15	BoE-BTL
Initial rent price mark-up, η	(9)	Estimated dist.	Zoopla, ONS
Rent price reduction			
Monthly probability	-	0.1057	Zoopla
Mean percentage reduction	-	1.6559	Zoopla
Std. dev. of percentage reduction	-	0.7855	Zoopla
Tenancy length			
Minimum	-	12	ARLA Survey
Maximum	-	24	ARLA Survey

Table 12: Housing decisions as a BTL investor

Parameter	Equation	Value	Source
Hard maximum LTV ratio			
First-time buyers	(12)	$\Gamma_{FTB} = 0.9$	MoneyFacts
Home movers	(12)	$\Gamma_{HM} = 0.9$	MoneyFacts
BTL investors	(12)	$\Gamma_{BTL} = 0.75$	MoneyFacts
Hard maximum LTI ratio			
First-time buyers	(13)	$\Phi_{FTB} = 5.4$	PSD
Home movers	(13)	$\Phi_{HM} = 5.6$	PSD
Hard maximum DSTI ratio for home buyers	(14)	$\Psi = 0.4$	PSD
Hard minimum ICR for BTL investors	(15)	$\Omega = 1.25$	BoE
Mortgage duration	-	25 years	PSD
Retirement age	-	65 years	Postulated
Elasticity of interest rate to credit	(17)	$\varphi = 1.33e^{-5}$	BoE, ONS
Bank initial interest rate	(17)	3.5%	BoE
Bank initial credit supply	(17)	£244.0	PSD, BoE-BTL
Exogenous policy rate	-	0.5%	BoE

Table 13: Bank and Central Bank parameters

Parameter	Equation	Value	Source
Bid-up parameter	(18)	$\chi = 1.0746$	Zoopla
Days under offer	(20)	3	Postulated

Table 14: Housing market parameters

B.2 Calibration: Method of simulated moments

After estimating from data as many parameters as possible, as well as manually calibrating those parameters for which a combination of an informed guess and a sensitivity analysis is sufficient, we need to calibrate the remaining 5 parameters. To do so, we use the method of simulated moments (Gilli & Winker, 2003; Franke, 2009; Franke & Westerhoff, 2012; Fabretti, 2013; Chen & Lux, 2018; Platt & Gebbie, 2018) to construct an objective function that measures the errors between a set of simulated moments and their corresponding data equivalents. Finally, by finding the minimum of this objective function over the hypercube of values explored for these parameters, we identify those values which minimise the distance between both sets of moments—simulated and estimated from data. Note that, in order to ensure being close enough to each of the moments selected, thus avoiding being too close in some and too far in others, we impose on each of them a bound of acceptability of $\pm 25\%$ of the target value.

The parameters calibrated with the method of simulated moments, as well as the specific values explored, are:

- Psychological cost of renting [see Eq. (4)]. Values explored: 0.0, 0.1, 0.2, 0.3, 0.4, 0.5.
- Sensitivity between rent or purchase [see Eq. (4)]. Values explored: 0.00001, 0.00003162, 0.0001, 0.0003162, 0.001, 0.003162, 0.01, 0.03162, 0.1.
- Buy-to-let probability adjustment [see Appendix A.2]. Values explored: 1.6, 1.62, 1.64, 1.66, 1.68, 1.7, 1.72, 1.74, 1.76, 1.78, 1.8.
- Sensitivity for BTL buy and sell decisions [see Eqs. (8) and (11)]. Values explored: 0.1, 0.3162, 1, 3.162, 10, 31.62, 100, 316.2, 1000.
- Weight of market vs segment house prices [see Eq. (21)]. Values explored: 0.1, 0.3, 0.5, 0.7, 0.9.

This amounts to 26730 parameter combinations. We run 10 simulations per parameter combination. Each simulation runs for 3500 time steps, out of which we will always discard the first 500 time steps, except for both sale and rental transactions, for which we directly avoid recording the first 1000 time steps.

The moments used by the method, as well as the specific values targeted, are shown in Table 15. As explained above, we ignore parameter combinations for which any of the targeted moments, when averaged over the 10 available simulations, lies outside of the corresponding bound of acceptability ($\pm 25\%$). For each of the remaining parameter combinations, we compute the value of the criterion function, defined as a weighted sum of average squared percentage distances to the target moments (average refers to the 10 simulations available). In the absence of any clear and universally accepted rule to select these weights, we chose to give extra weight (> 1) to the moments most relevant for the model output to match a realistic price behaviour. For example, we give higher weight to moments related to sale prices as compared to those related to rental prices. Finally, we simply choose the parameter combination corresponding to the minimum value of the criterion function.

Moment	Target Value	Simulation Value
House Price Index (HPI) mean	1.0	0.87
House Price Index (HPI) std. dev.	0.3424	0.2076
House Price Index (HPI) cycle period	201.0	186.63
Rental Price Index (RPI) mean	1.0	0.98
Share of households owning	0.65	0.62
Share of households renting	0.17	0.21
Share of households BTL investing	0.0753	0.0792
Rental yield mean	5.10	4.13
Residential mortgages spread mean	2.9985	2.9935

Table 15: Moments and their target values

Appendix C Online Appendix

This Online Appendix accompanies the paper “Heterogeneous Effects and Spillovers of Macroprudential Policy in an Agent-Based Model of the UK Housing Market”, by Carro, Hinterschweiger, Uluc & Farmer. It provides further detail on the estimation and calibration methodology of the parameters used in the model, as well as on the underlying data sources. The model code can be downloaded from <https://github.com/INET-Complexity/housing-model>.

Section C.1 describes user set parameters. Section C.2 provides information on the parameters that are estimated empirically. Section C.3 describes calibrated parameters which cannot be directly estimated from readily available data sources.

C.1 User set parameters

C.1.1 General model control

The parameters below (Table 16) are all to be set by the user for general model control.

Parameter	Value
Seed for random number generation	1
Number of time steps	10,000
Number of simulations	1
Number of households	10,000
Time step to start recording transactions	2,000
Rolling window for core indicator averages (in months)	6
Cumulative weight for events beyond 12 months	0.14

Table 16: User set parameters

C.1.2 Central Bank

These are all central bank control parameters determining the policies in place and thus used to run experiments. Their default values are estimated with data only if they are binding or influence the results in any way, which is the case only for the Central Bank base interest rate. The default value for this base rate is estimated as the actual Bank of England rate between 2009 and 2016. In particular, 0.5% was the base rate since 5th March 2009, when it was lowered from 1%, till 4th August 2016, when it was lowered to 0.25%. This parameter is kept at this default

value for all simulations. The rest of the parameters, shown in the below table, are identical to the ones that the bank sets in Table 13 in the main paper. They are therefore non-binding in the benchmark version of the model, and only bind when set to stricter values for the purpose of experiments, such as the ones undertaken in Section 5.

Parameter	Value
Hard maximum LTV ratio	
First-time buyers	0.9
Home movers	0.9
BTL investors	0.75
Hard maximum LTI ratio	
First-time buyers	5.4
Home movers	5.6
Hard maximum DSTI ratio for home buyers	0.4
Hard minimum ICR for BTL investors	1.25
Mortgage duration	25 years
Retirement age	65 years
Exogenous policy rate	0.5%

Table 17: Central bank parameters

C.2 Estimated parameters

These are all parameters which can be directly estimated empirically from available data sources.

C.2.1 General parameters

Number of households in the UK: Value 26,442,100. Used to compute the core indicators and the ratio of houses per household. Estimated with ONS 2011 Census data (Table H01UK) for the whole of the UK, households with at least one usual resident.

Number of dwellings in the UK: Value 22,626,000. Used to compute the ratio of houses per household. Estimated with ONS Table 101 (discontinued) for 2011 for the stock of dwellings in the UK. We use the sum of the number of owner-occupied and privately rented dwellings.

C.2.2 Household parameters

Age distribution: Density of households in each of 8 age bins between 15 and 95 years old. Estimated with WAS wave 3 household data.

Joint distribution of age and gross income: Household Representative Person's age bins and logarithmic income bins with their probabilities. Estimated with WAS wave 3 household data for total gross non-rent income (total gross income minus gross rent income).

Joint distribution of net financial wealth and gross income: Logarithmic gross income bins and logarithmic net financial wealth bins with their probabilities. Estimated with WAS wave 3 household data for total gross non-rent income (total gross income minus gross rent income) and net financial wealth. While WAS data includes also liquid wealth and gross wealth, the definition of net wealth is the closest to what we call financial wealth in our model. Liquid wealth is too restrictive, not including a number of types of wealth which can be easily transformed into liquid wealth and later used for house purchases. Gross wealth, on the contrary, would count as usable for house down-payments categories of wealth which would not be readily available for it.

C.2.3 Household behaviour parameters

Essential consumption fraction: Value 0.66. Fraction of Government income support (Job-seeker's allowance) necessarily spent a month as essential consumption. Estimated with Living Costs and Food Survey (LCFS) data for 2011 (collected by ONS). In particular, data for households with monthly incomes between £400 and £480 is used.

Maximum consumption fraction: Value 0.17. Maximum monthly desired consumption as a fraction of the household's gross annual income. Estimated with Living Costs and Food Survey (LCFS) data for 2011 (collected by ONS). In particular, we use the consumption fraction (monthly non-housing consumption divided by gross annual income, for incomes above £5900 for consistency with our income bins) corresponding to the 99th percentile.

Expectation formation: Slope (0.44) and intercept (-0.007) for the formation of expectations about the future behaviour of prices based on their past behaviour. Estimated with BoE's NMG Survey and Land Registry data for 2014/2018. House price growth is calculated over two years i.e. growth between the most recent quarter and the quarter two years before. This is a postulated parameter. Both a pre- and a post-calibration analysis show that this parameter has a fairly clear and contained effect mostly restricted to the period of the cycles. This analysis shows that

values of 1, 2 and 3 of this parameter can lead to price cycles coherent with OECD data for the UK. As a design decision the parameter value was set to the intermediate value of 2.

Desired purchase expenditure: Scale (42.9036); exponent (0.7892) to which the annual gross employment income of the household is raised when computing its budget for buying a house; and mean (-0.0177) and standard deviation (0.4104) of the normal noise used to create a log-normal variate, which is then used as a multiplicative noise when computing the desired budget. Empirically estimated using PSD and Land Registry data for 2011. PSD data was trimmed.

Desired down-payment: Scale (10.35) and shape (0.898) parameters for the log-normal distribution of desired down-payments by first-time-buyers. Empirically estimated with PSD data for 2011. Scale (11.15) and shape (0.958) parameters for the log-normal distribution of desired down-payments by home movers. Empirically estimated with PSD data for 2011. Average down-payment (0.34), as a percentage of house price for BTL investors and standard deviation (0.15) of the noise for down-payments by BTL investors. Empirically estimated with BoE - BTL data for 2014 (earliest available data point).

Holding period of owner-occupied houses: Value 17. Average period, in years, for which owner-occupiers hold their houses. $1/(12 \cdot \text{Hold Period})$ is the monthly probability that an owner-occupier household would sell its house. Empirically estimated from English Housing Survey (EHS) data for 2011.

Distribution of sale price mark-up: Initial sale price mark-up over average price of same quality houses, bins with their probabilities. Empirically estimated with Zoopla data (raw collated listings). In particular, final sale price, final sale date and initial posting date are used, combined with HPI evolution, to back-project an initial “real value” price for each property, which is then compared with the actual initial listing price to find an initial mark-up for each property.

Sale price reduction: Monthly probability of reducing the price of a house on the market is 0.0703. Empirical estimation with Zoopla data (raw collated listings) from 2003 to 2015. Mean (1.4532) and standard deviation (0.7070) of percentage reduction for prices of houses on the market. Empirical estimation with Zoopla data (raw daily listings) from 2003 to 2015.

C.2.4 Household behaviour: Buy-To-Let

Probability of receiving a BTL investor flag per income percentile bin: Income percentiles are found for each household in the data set using gross non-rental income, that is, gross total income minus gross rental income. Then, for each 1% bin, the fraction of BTL investors is found, using a non-zero gross rental income as flagging a BTL investor. Empirically estimated with WAS wave 3 household data for total gross non-rent income (total gross income minus gross rent income), and gross rental income.

Probability of BTL investor being driven by income, capital gains or mixed motives: The probability for rental-income-driven is 0.4927, for capital-gains-driven is 0.1458 and thus for mixed motives it is 0.3615. Probability of a BTL household having a rental-income-driven, capital-gains-driven or mixed strategy. For each household who receives the BTL flag, this probability decides the weights the household will assign to rental yield as opposed to capital gains. Empirically estimated with NMG survey data for 2014.⁴⁷

Weight on capital gains for BTL income-driven, capital-gains-driven and mixed strategies: Value 0.1 for rental-income-driven households, 0.9 for capital-gains-driven households and 0.5 for households with a mixed strategy. Sets the weight given to capital gains as opposed to rental yield by the three different types of BTL households. Design decision with robustness analysis both before and after full model calibration. These values are chosen as a representation of households highly driven by rental income, highly driven by capital gains or driven by an equal combination of both factors. Note that, in the case of households highly driven by capital gains or by rental income, they still take the alternative strategy into account. This avoids highly unreasonable behaviours, such as households sticking to their rental properties even if they believe prices would be completely tanking or households selling their rental properties if prices slightly fall even if rental income would be extremely high.

C.2.5 Household behaviour: Rent

Desired rental expenditure:

⁴⁷ Similar data is available from DCLG for 2010. However, this data refers to the fraction of houses and not the fraction of landlords. Finally, ARLA data is also available, but it includes other categories, such as saving for retirement, whose motivation is absent from our model.

- **Scale parameter of rental bid fraction:** Value 17.2166. Number of annual salaries (raised to the exponent of rental bid fraction) that the renter is willing to spend for renting a house. Empirically estimated using BoE's NMG Survey data for 2016.
- **Exponent of rental bid fraction:** Value 0.3464. Exponent to which the annual gross employment income of the household is raised when computing its desired budget for renting a house. Empirically estimated using BoE's NMG Survey data for 2016.

Rental length: Value 12 (minimum) and 24 (maximum). Minimum / maximum number of months of a rental contract. Rental contract lengths are drawn from a uniform discrete distribution between the minimum and maximum value for the tenancy length. Empirically estimated from literature using ARLA annual report (ARLA Members Survey of the Private Rented Sector) for 2013 Q4.

C.2.6 Household behaviour: Rent-out

Initial rental mark-up distribution: Logarithmic initial sale price mark-up over average price of same quality houses, bins with their probabilities. Empirically estimated with Zoopla data (raw collated listings) and ONS's RPI data (only for England, as full UK data is only available from 2015). In particular, final sale price, final sale date and initial posting date are used, combined with RPI evolution, to back-project an initial "real value" price for each property, which is then compared with the actual initial listing price to find an initial mark-up for each property.

Rent-out price reduction:

- **Probability of rental price reduction:** Value 0.1057. Monthly probability of reducing the rent price of an unrented house on the market. Empirical estimation with Zoopla data (raw collated listings) from 2003 to 2015.
- **Percentage reduction of rental price:** Value 1.6559 (mean) and 0.7855 (standard deviation). Mean / standard deviation of percentage reduction for rent prices of unrented houses on the market. Empirical estimation with Zoopla data (raw collated listings) from 2003 to 2015.

C.2.7 Housing market

Days under offer: Value 3. Time, in days, that a house remains under offer and thus the seller would consider bids for a possible bid-up. When an offer receives multiple bids in a single time step (a month), these bids are assumed to arrive randomly during the month. Due to the design of the market mechanism, which leads to a large numbers of bids per offer, the \log_{10} of the number of bids received is used here, to approximate a realistic number of bids. When a bid is followed by another bid within a time window of the “days under offer”, then the price is increased by multiplying it by the “bid-up multiple” (see below). As soon as a bid is followed by a time window of “days under offer” without another bid, the bid-up process is stopped. This parameter is considered a design decision, i.e. postulated. In any case, we have checked that the resulting distribution of the number of bid-ups is coherent with the one observed in Zoopla data. Furthermore, a robustness analysis done for a fixed set of parameters and trying different values for this specific parameter in the same order of magnitude shows only a small variation in terms of the distribution of the number of bid-ups.

Bid-up multiple: Value 1.0746. Size of a gazump or increase in price due to multiple bidders. Empirically estimated from Zoopla data (raw collated listings). The average price increase per price update is computed as $(\text{Price}/\text{Initial price})^{(1/\text{price updates})}$ for all houses with a positive price change.

Parameters for house price distribution: Value 12.1186 (scale) and 0.6414 (shape). Scale / shape parameter for the log-normal distribution of house prices. This estimated distribution of house prices is used to determine the initial price of each market quality segment. These initial sale prices are then used whenever households need to consider the sale price evolution of the different market quality segments, such as for computing the house price index HPI, which takes into account market composition at each time step. Empirically estimated with Land Registry Price Paid Data for 2011.

Parameters for rental price distribution: Value 6.2647 (scale) and 0.6353 (shape). Scale / shape parameter for the log-normal distribution of house rental prices. This estimated distribution is used to find the initial rental price of each market quality segment. These initial rental prices are then used whenever households need to consider the rental price evolution of the dif-

ferent market quality segments, such as for computing the rental price index RPI, which takes into account market composition at each time step. Empirically estimated with BoE's NMG Survey data for 2016.

Initial rental gross yield: Value 0.05. Initial profit margin for BTL investors at the beginning of the simulation. The effect of this value disappears soon after the model starts running and determining rental yields endogenously. Empirically estimated in the literature. In particular, we note two separate sources pointing at a 5% value. First, an analysis matching data from Zoopla and Land Registry shows yields in the area of 5 – 6.6%, depending on the type of property.⁴⁸ Second, ARLA's annual report for 2013 points at an average value of 5% for rental yields.

C.2.8 Bank

Mortgage duration: Value 25. Mortgage duration in years. Empirically estimated with PSD data for 2011 as the median mortgage duration in this dataset.

Retirement age: Value 65. Maximum age for a household to get a mortgage and maximum age for a non-BTL household to finish repaying its mortgages, i.e. akin to their retirement age. BTL mortgages have always full maturity, but they can only be approved before the household reaches this age limit. Non-BTL mortgages, on the contrary, will have their maturities reduced before this age limit in such a way that they are fully repaid by the time the household reaches this age limit.

Bank initial interest rate: Value 0.035. Private bank initial interest rate for mortgages. Note this is the total rate, not the private bank spread. The effect of this value disappears soon after the model starts running and setting interest rates endogenously. Empirical estimation with Bank of England Housing Core Indicators, overall spread on new mortgages plus central bank rate, averaged for all months in 2011.

Bank initial credit supply: Value 244.0. Bank's initial supply of credit per household per month. It is used as initial value of the previous supply of credit per household per month, which is then compared with the actual supply of credit per household per month in order to

⁴⁸ Bracke, P. (2015). *How much do investors pay for houses?* Staff Working Paper 549, Bank of England.

determine the interest rate movement for next month. Empirically estimated with PSD data for 2011 and BoE-BTL data for 2018. In particular, principals advanced to first-time buyers and home movers are assessed from PSD data for 2011 while principals advanced for BTL investment are assessed from BoE-BTL data for 2018.

Elasticity of interest rate to credit: Value $1.33e^{-5}$. Rate of change of the interest rate in response to a change in the demand/supply for credit per household per month (in points per pound per household per month). Empirically estimated with a combination of data sources. Since we are only interested in setting the order of magnitude of rate movements as a response to credit demand changes, we perform here a very rough estimate. In particular, the value of this parameter is the average of the absolute value of month to month interest rate differences (from BoE core indicators) divided by the average of the absolute value of month to month amount of credit per household differences (from BoE stats table A5.3 VTUZ and ONS number of households). All data sources are trimmed to the common period, 1996 – 2017.

Hard maximum LTV ratio: Value 0.9 (both for first-time buyers and home movers). Internal bank policy, hard maximum LTV ratio for first-time buyers and home movers. Empirically estimated with MoneyFacts data for 2011 (which contains information on the products offered). In particular, we use the mode of the 10 largest UK mortgage lenders which report MoneyFacts data.

Hard maximum LTV ratio for BTL investors: Value 0.75. Internal bank policy, hard maximum LTV ratio for BTL investors. Empirically estimated with MoneyFacts data for 2011 (which contains information on the products offered). In particular, we use the mode of the 10 largest UK mortgage lenders which report MoneyFacts data.

Hard maximum LTI ratio: Value 5.4 (first-time buyers) and 5.6 (home movers). Internal bank policy, hard maximum LTI for first-time buyers / home movers. Empirically estimated with PSD data for 2011. In particular, we use the 99th percentile.

Hard maximum DSTI ratio for home buyers: Value 0.4. Internal bank policy (debt-service-to-income or debt-service ratio), hard maximum fraction of household's income to be spent on mortgage repayments (for both first-time buyers and home movers). Empirically estimated with PSD data for 2011. In particular, we use the 99th percentile.

Hard minimum ICR for BTL investors: Value 1.25. Internal bank policy (interest coverage ratio), hard minimum ratio between (expected) annual rental income and annual interest payments. Industry standard described as 125% in PRA Supervisory Statement SS13/16 (paragraph 2.7).

C.2.9 Government parameters

Income tax personal allowance: Tax-free personal allowance for income tax if income does not go above the personal allowance income limit (see next parameter). Estimated with HMRC data on income tax personal allowances and reliefs for the tax year 2011-2012. In particular, we use the personal allowance for a single person, £7475. We have, however, checked Wealth and Assets Survey data for 2011 in order to confirm that a single allowance does lead to net income results closer to data, as opposed to using a double allowance which would rather correspond to a household with two tax payers.

Income limit for income tax personal allowance: Income limit above which the untaxed personal allowance (see previous parameter) starts to decrease at a rate of £1 for every £2 the gross income goes above this limit. Estimated with HMRC data on income tax personal allowances and reliefs for the tax year 2011-2012. Its value is £100000.

Minimum government income support: Income support (Jobseeker's allowance) for a couple, both over 18 years old, in 2011, which amounts to £445.80 a month. Estimated with DWP data on government support and allowances for the year 2012.

Income tax: Bands of income (lower thresholds after personal allowance) and their corresponding tax rates. Estimated with HMRC data on income tax for the tax year 2011-2012. These tax rates are applied to all income in the model, thus including also rental income. Finally, we have also implemented finance costs relieve for BTL investors, exempting them from paying tax on their interest costs. Note that, while this finance cost relief was progressively removed between 2017 and 2020, it was available during our target year for estimation, 2011.

National insurance contributions: Bands of income and their corresponding tax/contribution rates. Estimated with HMRC data on national insurance contributions for a class 1 employee for the tax year 2011-2012. Note that the bands used already include any untaxed allowance.

C.3 Calibrated parameters

These are parameters which cannot be directly estimated from readily available data sources, whose effects are widespread across the model and for which it is unfeasible to postulate plausible values.

C.3.1 Household behaviour: Rent vs purchase decision

Psychological cost of renting: Annual psychological cost of renting, as a percentage increase over the actual cost of renting. The values tested in the calibration process went from 0.0, so as to not impose the mechanism, up to 0.5 in steps of 0.1. Pre-calibration analysis pointed at values above 0.5 as leading to too low rental prices, since only households unable to buy would decide to bid in the rental market. The value selected by the calibration process was 0.4.

Sensitivity between rent or purchase: Sensitivity parameter for the decision between buying and renting. We have tested 9 logarithmically arranged values between 0.00001 and 0.1, thus exploring 5 orders of magnitude, 2 below order $O(0)$ and 2 above it, plus intermediate points. In this case, since the variable multiplied by the sensitivity parameter in the exponential has order of magnitude $O(3)$, then the central order of magnitude for this sensitivity parameter will be $O(-3)$, that is, 0.001. The value selected by the calibration process was 0.001.

C.3.2 Household behaviour: Buy-To-Let

Buy-to-let probability adjustment: Adjustment or multiplier factor for the probability of any household receiving the BTL flag, thus allowing it to invest in BTL properties. This, thus, modifies the probability of receiving this flag for all income percentile bins, as estimated from the Wealth and Assets Survey. The purpose of this parameter is to increase these probabilities, such that the actual number of BTL investors matches the real one, also estimated from the Wealth and Assets Survey to be 0.075262. In this way, this takes into account the fact that what we are setting is the (unobservable) tendency to become a BTL investor, but the data contains the actual number of investors. For different reasons, such as having insufficient financial wealth or a lower income, a number of households in the model with the BTL flag will be unable to buy investment properties. Thus, we need a larger fraction of households with the flag so as to

match the actual fraction of investors. Since pre-calibration analysis pointed at values below 1.6 as leading to too few actual investors and values above 1.8 as leading to too many of them, we have tested 11 values between and including these two extremes during the calibration process. Finally, the value selected was 1.76.

Sensitivity for BTL buy and sell decisions: Sensitivity parameter for the decisions to buy and to sell BTL investment properties, as a function of their expected yield. As with the previous sensitivity parameter, we have tested 9 logarithmically arranged values, this time between 0.1 and 1000, thus exploring 5 orders of magnitude, 2 below order $O(0)$ and 2 above it, plus intermediate points. In this case, since the variable multiplied by the sensitivity parameter in the exponential has order of magnitude $O(-1)$, then the central order of magnitude for this sensitivity parameter will be $O(1)$, that is, 10. The value selected by the calibration process was 100.

C.3.3 Housing market

Weight on market vs segment house prices: This is the weight given by households to the general market behaviour of prices, as measured by the HPI, as opposed to the price behaviour of the specific market quality segment under consideration. Thus, it is used to compute the exponential moving average transaction price for the different quality segments. A value of 0.0 would mean households focusing only on the price evolution of the whole market, completely disregarding the price evolution of the specific quality segment of interest, while a value of 1.0 would mean households focusing only on the quality-segment-specific price evolution, disregarding the behaviour of the market as a whole. In addition to these two extreme cases, we have explored the whole range of this parameter in steps of 0.1. The value selected by the calibration process was 0.5.