



BANK OF ENGLAND

Staff Working Paper No. 619

Macroprudential policy in an agent-based model of the UK housing market

Rafa Baptista, J Doyne Farmer, Marc Hinterschweiger,
Katie Low, Daniel Tang and Arzu Uluc

October 2016

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Authority Board.



BANK OF ENGLAND

Staff Working Paper No. 619

Macroprudential policy in an agent-based model of the UK housing market

Rafa Baptista,⁽¹⁾ J Doyne Farmer,⁽²⁾ Marc Hinterschweiger,⁽³⁾ Katie Low,⁽⁴⁾ Daniel Tang⁽⁵⁾ and Arzu Uluc⁽⁶⁾

Abstract

This paper develops an agent-based model of the UK housing market to study the impact of macroprudential policies on key housing market indicators. This approach enables us to tackle the heterogeneity in this market by modelling the individual behaviour and interactions of first-time buyers, home owners, buy-to-let investors, and renters from the bottom up, and observe the resulting aggregate dynamics in the property and credit markets. The model is calibrated using a large selection of micro-data, mostly from household surveys and housing market data sources. We perform a series of comparative statics exercises to investigate the impact of the size of the rental/buy-to-let sector and different types of buy-to-let investors on housing booms and busts. The results suggest that an increase in the size of the buy-to-let sector may amplify house price cycles and increase house price volatility. Furthermore, in order to illustrate the effects of macroprudential policies on several housing market indicators, we implement a loan-to-income portfolio limit. We find that this policy attenuates the house price cycle.

Key words: Agent-based model, housing market, macroprudential policy, buy-to-let sector.

JEL classification: D1, D31, E58, R2, R21, R31.

(1) Institute for New Economic Thinking (INET) at the Oxford Martin School, and Department of Computer Science, University of Oxford. Email: rafa.baptista@some.ox.ac.uk

(2) Institute for New Economic Thinking (INET) at the Oxford Martin School; Mathematical Institute, University of Oxford; and Santa Fe Institute, New Mexico. Email: doyne.farmer@inet.ox.ac.uk

(3) Bank of England. Email: marc.hinterschweiger@bankofengland.co.uk

(4) Bank of England. Email: katie.low@bankofengland.co.uk

(5) Institute for New Economic Thinking (INET) at the Oxford Martin School, and Mathematical Institute, University of Oxford. Email: daniel.tang@inet.ox.ac.uk

(6) Bank of England. Email: arzu.uluc@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. The authors would like to thank John Muellbauer, David Miles, Andy Haldane, Nicola Anderson, Sujit Kapadia, Gavin Wallis, David Pugh, Marco Galbiati, Tracy Wheeler, Philippe Bracke, Sagar Shah, David Bradnum, Ian Billett, Joshua Wignall, participants of the Family Finance Surveys User Conference 2016, participants of a seminar at the Central Bank of Ireland, as well as participants of a seminar at the Bank of England for their valuable comments and support. Any errors and omissions, of course, remain the fault of the authors.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

1 Introduction

Housing is the largest asset class in the world and an important sector of the economy. Developments in the housing market can be significant drivers of economic dynamics: house price growth, credit growth, the safety of bank balance sheets, and household indebtedness, among others, are closely intertwined.

Several features of the housing market contribute to its prominence. First, housing finance is often leveraged, aggravating housing and credit boom-bust cycles. Second, housing is a relatively illiquid asset, which can cause difficulties for sellers, especially in distressed circumstances. Third, housing is local, making adjustments in supply and demand sluggish. Fourth, houses are a heterogeneous good, making it difficult to model using aggregate methods. Fifth but not least, it involves a complicated interaction between several different kinds of agents, including renters, first-time buyers, home-movers and buy-to-let investors.

From a modelling perspective, these features are often difficult to capture properly with the standard tools available to economists. Generally, researchers choose convenient short-cuts that represent a good approximation in most circumstances, but fail strikingly in some, especially in policy design. For example, mainstream economic models such as dynamic stochastic general equilibrium (DSGE) models, which have a rational representative agent at their core, often rely on overly simplified assumptions that can cause models to differ profoundly from reality in some cases and lead policy-makers astray (Farmer and Foley, 2009).

Agent-based models (ABMs) can overcome these shortcomings to a certain degree. The ABM approach enables the modelling of different actors in the housing market in detail. It allows for non-linear dynamics, such as housing booms and busts, which come naturally from the interactions of buyers and sellers in the market. Furthermore, it permits an evaluation of the impact of policies which target a certain segment of the market, such as the Bank of England's loan-to-income (LTI) flow limit that came into force in October 2014¹. That policy requires mortgage lenders to

¹ See <http://www.bankofengland.co.uk/prd/Documents/publications/ps/2014/ps914.pdf>

constrain the proportion of new lending at LTI ratios at or above 4.5 to no more than 15% of the total number of new mortgage loans. Such heterogeneity (i.e. some agents receiving mortgages at an LTI ratio above 4.5, while others do not) is straightforward to capture in an ABM.

ABMs are not without drawbacks. Behavioural rules are often arbitrary, calibration is difficult, and causality is frequently obscure. However, a model of the housing market that can account for housing's special characteristics in an ABM framework can be useful to policy-makers.

This paper describes an agent-based model of the UK housing market. The aim is to study the possible effects of macroprudential policies, including on nine key housing market indicators monitored by the Bank of England's Financial Policy Committee (FPC).

In the model, there are three types of agents: (i) households, (ii) a bank (mortgage lender), and (iii) a central bank. Households are the main agents in the model. There are four types of households: renters, first-time buyers, owner occupiers, and buy-to-let (BTL) investors. All households differ along certain characteristics, such as their age, income, and bank balances. They buy, sell, lease and rent houses. When buying houses, households can take out mortgages from a bank which represents the mortgage lending sector in the aggregate. The mortgage market is subject to regulation by a central bank that sets macroprudential policies.

The model is calibrated against a large set of micro data, mostly from household surveys and housing market data sources, e.g. the Financial Conduct Authority's (FCA) loan-level Product Sales Data (PSD)², Land Registry transactions data, and WhenFresh/Zoopla data on rental listings. The calibration of the model proceeds in two steps: a micro-calibration that fine-tunes household's individual characteristics and behaviour (e.g. the joint income and age distribution); and a macro-calibration that ensures consistency with economic aggregates (e.g. the FPC's nine housing core indicators). The model does not attempt to represent the state of the UK housing market at a specific point in time; in particular, it is not its objective to produce forecasts. Rather, our approach is intended to model the baseline or long run dynamics of the market, focusing on the different

² The FCA Product Sales Data include regulated mortgage contracts only. See <http://www.fca.org.uk/firms/systems-reporting/product-sales-data/> for officially published high level data.

agents that create these dynamics.

The calibrated model is then used to perform a series of comparative statics exercises. First, we investigate the impact of the size of the rental/BTL sector and of different types of BTL investors on housing booms and busts. Then, in order to illustrate the effects of macroprudential policies on several housing market indicators, we implement a loan-to-income portfolio limit on borrowers' ability to receive mortgages, allowing a share of new mortgages to be unconstrained by this limit. We find that this policy attenuates the house price and credit (mortgage) cycle.

This paper contributes to the literature by expanding existing housing ABMs along different dimensions, especially by introducing realistic life-cycle dynamics including death and inheritance, a rental and BTL sector, and a double-auction market mechanism. 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.

This paper is structured as follows. Section 2 describes the literature on housing and ABMs. Section 3 sets out the model. Section 4 explains the calibration of the model, while Section 5 describes how the model can be validated. Section 6 presents the analysis of the BTL market and a loan-to-income policy experiment. Section 7 concludes.

2 Literature review

Although most macroeconomic models of the housing market and housing policies are DSGE models, there has been an increasing interest in developing housing market models using an agent-based approach in recent years.

The seminal paper by Iacoviello (2005) introduces a housing collateral constraint and nominal debt contracts in a standard DSGE model. The starting point is a new-Keynesian DSGE model with a financial accelerator following Bernanke et al. (1999). In the model, a monetary policy

shock leads to a fall in house prices, reducing the value of collateral and hence borrowers' access to credit. This reinforces the standard transmission channels of monetary policy.

Several DSGE models with housing focus on the effect of policies that are, unlike monetary policy, explicitly targeted at the housing sector. Lambertini et al. (2013) find that a countercyclical loan-to-value (LTV) policy that stabilizes credit relative to GDP is optimal with regards to borrowers. Mendicino and Punzi (2014) examine the interaction of house prices, the current account, and debt in a two-country setting, determining an optimal policy mix between an interest-rate response and an LTV policy. Kannan et al. (2009) highlight the importance of the source of a house price boom (such as from a productivity or demand shock) when designing the optimal policy response. Gelain and Bank (2011) and Gelain et al. (2012) examine the effects of different macroprudential instruments such as a loan-to-value, loan-to-income, or leverage ratio on house prices and credit. In a case study on Hong Kong, Funke and Paetz (2012) find that LTV policies that are designed to reduce house price growth can mitigate the effects of house price cycles on the real economy.

All these DSGE models have in common that the only purpose of housing is to serve as collateral for borrowing, which is driving the dynamics in these models. Furthermore, the housing sector is homogeneous in these models, i.e. there are no different segments such as renters, owner-occupiers or investors with distinct motives. Last but not least, even when papers study constraints that only bind occasionally, the degree of 'bindingness' is often declared to not have an effect large enough to merit a detailed analysis. For instance, the papers by Iacoviello and Neri (2010) as well as Lambertini et al. (2013) solve their models by linearising the equilibrium conditions of the model around a steady state with a binding borrowing constraint.

This seems to contradict the experience of real-world policy-making: policy-makers are often precisely concerned about the subgroup of the population or the subset of product that are constrained. For example, housing policies are often designed to help potentially credit-constrained customers get mortgages (such as the UK government's 'Help to Buy: mortgage guarantee scheme'³)

³ See <https://www.helptobuy.gov.uk/mortgage-guarantee/how-does-it-work/>

or targeted at those types of mortgages that can exacerbate financial risks (such as the limit on high loan-to-income mortgages mentioned above). Rather than treating these cases as exceptions, ABM models are methodologically well suited to provide policy-makers with analysis that sheds light on precisely these cases.

Our approach is based on Axtell et al. (2014), preliminary results for which were presented in Geanakoplos et al. (2012), who made a detailed model of the housing market of Washington, DC. This work is novel in the literature as a wealth of micro data sources were used to calibrate behavioural equations, as opposed to postulating top-down, theoretical behavioural rules, as in most of the other models mentioned in this section. The main focus of their work was to demonstrate the causal relationship between leverage and the formation of a housing bubble. Their experiments showed that loan-to-value constraints were a major determinant of bubbles, more so than interest rates.

Our model has followed the same overall approach in the use of multiple sources of micro data to elicit behaviours. Several design choices have been derived from Axtell et al. (2014), including the representation of houses characterised by a quality parameter, the mortgage lending sector dynamics represented by a single bank enforcing lending constraints and a central bank imposing regulations, the method of determining the selling price and the downpayment, and the structure of the main simulation loop. Our model has added to this by using a more general double-auction market mechanism, and has a variety of new features, include a new consumption rule, life-cycle dynamics including death and inheritance, and a more empirically grounded treatment of renting and BTL investors.

A number of housing market ABMs have introduced a spatial component in their representations. Modelling households and/or houses with spatial attributes can be a highly desirable feature of an ABM, since it is necessary in order to capture location-dependent dynamics, such as regional bubbles or migration patterns. However, this approach greatly increases the complexity of the models and hence most spatial ABMs in the field listed below make use of a highly simplified

representation of the environment, often in the shape of small grids. Gilbert et al. (2009) modelled the interaction between buyers, realtors and sellers in an agent-based model of the English housing market, by making use of a grid representing a small city and focusing on the appearance of regionally clustered price variations and on the effect on the area of influence of realtors. Ge (2014) develops a spatial model of the housing market to study the formation of real estate bubbles. McMahon et al. (2009) perform a similar exercise, introducing shocks in the interest rate into a spatial grid of properties, and modelling foreclosure. Pangallo et al. (2016) employ a mathematical approach using a continuous Euclidean space in order to reproduce segregation patterns and the influence of income inequality on their emergence. Devisch et al. (2009) model in great detail the agents' housing choice and price negotiation processes by postulating a comprehensive utility function including a spatial component, though they focus less on the macroeconomic implications that result from such choices and from the interactions among the agents.

There are also two other papers analysing the housing market in an agent-based context. Kouwenberg and Zwinkels (2015) replicate historic boom-bust cycles in the US housing market by postulating a mix of fundamentalist and trend-following investors who buy and sell properties; the endogenous cycles arise from the heterogeneity in the calculation of price expectations. A model by Erlingsson et al. (2014) focuses on the macroeconomic implications of an easy access to credit and, by using different types of agents including households, firms, banks, funds and government, they find that loose credit conditions inflate house prices and create more unstable economies.⁴

3 Model description

There are three types of agents in the model: (i) households, (ii) a bank (mortgage lender), and (iii) a central bank. Households are the main agents in the model. Each period, some households

⁴ A related literature that analyses the housing market in an agent-based context concentrates on land use and urban development (for example, Diappi (2008); Jordan et al. (2010); Magliocca et al. (2011)). However, the focus of this strand is quite different from our own, which is more concerned about the economic rather than social or geographic aspects of the housing market.

are born, some households die and others age. All households receive income, spend on non-housing consumption, make decisions regarding their housing choice, whether owning or renting, pay housing expenses, and save the rest of their income. In addition, some households have a randomly allocated “BTL gene” that can lead them to trade (buy and sell) additional properties which can be let out to renters. When buying houses, households can take out mortgages from a bank which represents the mortgage lending sector in the aggregate. The mortgage market is subject to regulation by a central bank that sets macro-prudential policies. Figure 1 shows a simplified representation of the interactions between different agents in the housing market.

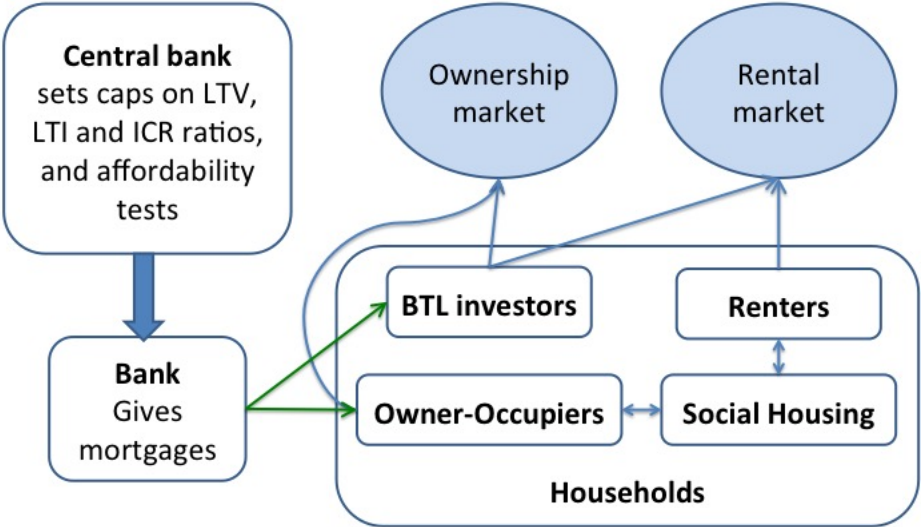


Figure 1: Model diagram showing the interactions between the model components

In this section behavioural rules and decisions of agents are presented in detail. The reader is referred to the Appendix B for an overview of the meaning of all variables and parameters contained in the model developed below.

3.1 Overview of the simulation step

Section 3.1. provides a brief overview of the different steps involved in simulating the model. The following sections go into more detail for each of the steps. Please refer to the Appendix B for a list of sources motivating the choice of parameters.

Once the model has been initialised (see the Appendix A.1), it proceeds in simulation steps of one month, where the following components interact in each iteration:

1. **Demographics:** birth, death and inheritance. New households are born at a constant birth rate, and assigned an age and income percentile from an appropriate distribution. All other households age, and some die according to an age-dependent distribution. When a household dies, another randomly selected household inherits its properties and wealth (see Section 3.2).
2. **Construction (optional):** Houses are characterized by a quality parameter, which is a proxy for size, location and conditions. If there are not enough houses in the system, new houses are added by a construction sector. The housing stock in the model is kept roughly constant. The target ratio of houses to households is 82%. If there is a shortfall, new houses are put on the market at a relevant price.
3. **Households:**
 - (a) Households receive income according to a distribution as a function of income percentile (assumed fixed over a lifetime) and age. Investor households collect rent if applicable. They pay taxes and spend on non-housing consumption (see Section 3.3).
 - (b) Bankruptcy (optional): it is possible at this point that a household cannot afford its mortgage payments or rent, therefore becoming bankrupt. Because the model does not seek to capture bankruptcy dynamics (such as default and foreclosure), we ensure that payments are always made by artificially injecting as much cash as necessary to the bankrupt households.

(c) Households make their housing decisions, according to their current status:

- If in social housing or if their renting contract just expired, they decide whether to buy or rent a new house.
- If renting, they continue to pay rent (tenants never decide to leave a renting contract before it expires).
- If owner-occupying a house, they decide whether to sell it and either buy a new one or switch to renting.
- If having BTL gene, they additionally decide whether to buy any new properties, and for each owned property, they decide whether to sell it.

(d) Households that have decided to buy or rent a new house place their bids on the relevant market (ownership or renting). Households that have decided to sell or rent out any properties also place their offers in the respective markets.

(e) Renters pay rent, and owners with a mortgage make their mortgage payments.

4. **Market clearing:** the ownership market is cleared first, followed by the rental market.

(a) In the ownership market, offers from sellers and buyers are matched in a number of rounds. Unsuccessful selling bids that were not matched to any buyers are “sent back” to the sellers, who can update or remove them, and possibly put their investment properties on the rental market instead. Offers that were matched to multiple buyers have the option of bidding up their demanded price.

(b) The rental market is cleared afterwards in the same way.

5. **Bank’s mortgage interest rate:** the bank (which represents the whole lending sector) calculates the mortgage interest rate, depending on its credit supply target (which is fixed) and the credit it supplied in the last month.

6. **Core indicators and Central Bank policy update:** finally, the core indicators are computed

and, if applicable, the central bank updates its regulations according to the new values of the core indicators.

Note: in the following sections, lower case Greek letters denote tunable parameters. For the sake of clarity, given the high number of parameters, we have named all parameters $\alpha, \beta...$ in every equation. The specific values of the parameters can be found in the Appendix B, by referring to the corresponding Equation number.

3.2 Demographics

3.2.1 Birth

Each month a number of households is born according to a fixed birth rate. Upon birth, a household is endowed with an age⁵, income percentile (which is assumed to remain constant through their life) and starting wealth. A fixed proportion (8%) of households that are above the 50th percentile of income are given a BTL ‘gene’ which gives them the desire to enter the BTL market (we refer to them as *investors*). All households are born into “social housing”⁶ but immediately make a decision of either renting or buying. Social housing represents homelessness, living with parents while looking for a house or living in temporary accommodation (e.g. hotel) while between houses.

⁵ The age of a household represents the age of the ‘household reference person’ (HRP) (a concept that exists in many household surveys). The birth of households can happen either by children leaving their parents’ home or by households splitting due to a divorce. As the HRP only accounts for leaving home but not for new households created by divorce, our calibration effectively only accounts for the former. ⁶ Agents never choose to be in “social housing”, but are put there if they fail to secure any other form of housing at a given time. If they find themselves in social housing they will always consider renting or buying and will bid on the appropriate market. When in social housing, no rental payments are deducted from income, this is a very simple form of housing benefit.

3.2.2 Mortality

In the model, the ‘death’ of a household implies that all its occupants have died⁷. We use a probability density function (pdf) of household death with respect to age and a pdf of the household ages. Integrating the two distributions together yields the overall death rate. In order to achieve a constant population in the model, this has been multiplied by a constant factor to set it equal to the birth rate. Details of the pdfs involved are given in the Appendix B.

3.2.3 Inheritance

Upon death of a household, all of its financial and housing wealth is given to another, randomly chosen, household. If the deceased household had any houses on any market, these are taken off the market. Outstanding mortgages are written off. Any tenants living in the houses are evicted. If the household was renting, the rental contract is terminated.

Upon being given a house, renters and the socially housed will immediately move into it, an owner-occupier will immediately sell it and a BTL investor will decide whether to sell it or add it to its portfolio.

3.3 Income and financial wealth processes

3.3.1 Income, housing consumption and tax

As a household ages, its income percentile remains fixed but its actual income changes 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. Income is bounded by a lower limit of £5,900 (the current level of income support for a married couple), and consists of employment income and collected rent (if the household is an investor and is renting out properties), minus mortgage or rent payments, National Insurance and income tax.

⁷ In reality, ‘death’ of a household can also be caused by marriage (two households fusing into one) or by friends moving in together.

3.3.2 Non-housing consumption and financial wealth

In order to justify the implementation of household consumption, we must consider its effect on the whole system. Because the model does not include macro-economic mechanics such as economic growth or unemployment, it is not necessary to model the effect of consumption on those indicators. Instead, the only effect of consumption in the model is to determine household *wealth*, which in turn affects the downpayments and available mortgages.

Therefore, we need to ensure that household consumption leads to an accurate wealth distribution for the population, which we obtained from the Wealth and Assets Survey. We can link household consumption to the wealth distribution as follows. First, we subtract a fixed, essential consumption, set at the lowest level of income support for a married couple (5900 a month). Then, for each household, we define the bank balance w they should have at the end of the month for the whole population to match the distribution obtained from the household survey; we call this quantity ‘desired bank balance’. The relationship between wealth and income in the Wealth and Assets Survey was clearly log-normal, and therefore, we define our desired bank balance w as

$$\ln(w) = \alpha + \beta \ln(y) + \varepsilon, \quad (1)$$

where y is income and ε is a Gaussian noise term introducing heterogeneity. However, if we were to simply set the household consumption so that the bank balance at the end of the month equals the desired bank balance w , we could find situations whereby households with higher incomes consume less. We can avoid this situation by forcing consumption C to be a fraction α of the difference between the current bank balance b and the desired bank balance w , according to

$$C = \max(\alpha(b - w), 0). \quad (2)$$

This equation ensures that the wealth of every household will relax exponentially towards the expected value w , while also ensuring that households with higher incomes consume more.

3.4 Housing decisions

At each time-step (every month), the following behavioural rules are applied to every household in the model, representing their housing decisions. The decisions available depend on their current status: whether they are (i) in social housing (i.e. seeking a home), (ii) renting, or (iii) owning-occupying. Additionally, if the household is an investor household, there will be extra options (iv) available.

3.4.1 Decisions if in social housing

All households needing a new home are represented as being in social housing (including new entrants, renters whose contract ended and owners who just sold their home). These households must decide whether to buy or rent a new house. In order to choose, they decide how much they want to spend given their available mortgage, find out what house quality they can afford, estimate and compare the costs of renting and owning, make the choice between the two and place the corresponding bid on the market. In detail:

1. The desired expenditure, $p_{desired}$, is given by

$$p_{desired} = \frac{\alpha y \exp(\varepsilon)}{1 - \beta g}, \quad (3)$$

where y is income, g is the expected monthly house price appreciation defined in Equation (4) below, and ε is noise. This equation is identical to the one used by Axtell et al. (2014) with different parameter values. This equation can be understood as setting the desired expenditure so that the long-term cost of the house (which takes into account the expected house capital appreciation) is a noisy fraction of income.

The actual desired expenditure p is set as close as possible to $p_{desired}$, but can be limited by the bank's mortgage available to the household.

2. The household finds out the house quality Q they can afford for a price p given the current

market conditions. This house quality will be used as a basis for the comparison between the cost of owning and renting.

3. The household will estimate the costs of the two options. The annual cost of renting is $r_Q(1 + \tau)$: the annual rent r_Q for a house of quality Q at the current market conditions, multiplied by a factor of $(1 + \tau)$ which represents the psychological cost of renting. On the other hand, the annual cost of owning is $12(m - pg)$, in terms of the monthly mortgage payment m minus the house appreciation pg , where p is used assuming that the household will buy a house worth its desired expenditure, and g is the household's estimation of the house price index change estimated for the current month. The estimate g is made by using the quarterly HPI of last year compared to this year's, according to

$$g_t = \alpha \left(\frac{h_{t-1} + h_{t-2} + h_{t-3}}{h_{t-13} + h_{t-14} + h_{t-15}} - 1 \right), \quad (4)$$

where α is a constant parameter and h_i is the monthly house price index for the month i , where h_t corresponds to the current month.⁸

4. The discrete choice of buying vs. renting as a function of the comparison between the cost of each option has been modelled by using the logistic function $\sigma(\beta x) = 1/1 + e^{-\beta x}$ (sometimes called standard sigmoid function), where the variable x is the difference between the two costs, and β is the shape parameter for the logistic function, which acts as a sensitivity parameter (the higher β , the more deterministic the decision is with respect to the variable x). The logistic function is a common choice whenever a continuous variable determines the probability of a discrete outcome with a clear threshold (it is often used as an 'activation function' in biology and computer science). Combining the definitions above we arrive at

⁸ See the Appendix A.3 for house price index calculation. This is based on empirical evidence that suggests last year's house price growth is an important determinant for forming house price expectations. Although this way of forming expectations may be seen as naive, it describes the behaviour observed in the data fairly well.

the behavioural rule of deciding to buy as opposed to rent, given by

$$\begin{aligned} P(\text{buy}) &= \sigma(\beta \cdot [\text{RentingCost} - \text{BuyingCost}]) \\ &= \sigma(\beta \cdot [r_Q(1 + \tau) - 12(m - pg)]). \end{aligned} \tag{5}$$

5. Finally, the household will place its bid on the market. If it decided to buy, it will bid its desired expenditure p on the ownership market, and if it decided to rent, it will bid a fixed fraction of its income (set at 33%) on the rental market.

Downpayment. Once the decision to buy is made, households make their purchases either outright or by obtaining a mortgage. If the financial wealth of the household is greater than the price of the house by a factor of 2, they will pay for the house in cash.⁹ Otherwise, they will request a mortgage and choose a downpayment. The minimum downpayment is determined by the mortgage conditions set by the bank (in Section 3.6); however, the household may choose to make a larger downpayment. A relationship between downpayment and income has been obtained from PSD; see the Appendix for more details.

3.4.2 Decisions as a renter

Households currently on a renting contract will continue to rent until the contract expires (unless they get evicted due to a deceased landlord, or upon inheriting a property). Therefore, a household will never choose to terminate their renting agreement. When the contract ends, they are put in social housing and hence the rules applying to them are those in Section 3.4.1.

⁹ This ensures a roughly correct proportion of cash buyers, calibrated against WhenFresh/Zoopla data.

3.4.3 Decisions as a homeowner

Households sell their owner-occupied property on average every 11 years¹⁰ due to exogenous reasons not addressed in the model, such as starting a family or divorce. The probability that a household will sell its home in a given month is

$$P(\text{sell}) = \frac{1}{12} \max \left(\frac{1}{11} \left[1 + \alpha(\bar{n}_h - n_h) + \beta(\bar{i} - i) \right], 0 \right), \quad (6)$$

where n_h is the number of houses per capita currently on the market, i is the interest rate as a percentage, \bar{n}_h, \bar{i} are the exponential moving averages of the corresponding quantities, and α and β are parameters. The term on n_h is introduced in order to prevent unrealistic build-up of housing stock on the market and the term on i prevents unrealistic fluctuations of the interest rate. Since the averages \bar{i} and \bar{n}_h are calculated during the simulation, it can be seen that the long-term selling probability converges to $1/11$, that is, selling once every 11 years.

If the household decides to sell its home according to the equation above, it will offer it on the market at a price p_s given by

$$\ln p_s = \alpha + \ln(\bar{p}) - \beta \ln(\zeta(1 + \bar{f})) + \varepsilon, \quad (7)$$

where \bar{p} is the average sold-price of houses of this quality, \bar{f} is the average number of days on the market for all house qualities, ε is Gaussian noise, and α, β and ζ are parameters. This is the same equation used by Axtell et al. (2014) with different parameter values. Their choice was motivated by observed exponential relationships found in transaction data for Washington, DC between the selling price, the average price of houses of similar quality or ‘similar houses’¹¹, and

¹⁰ Average sale rate based on English Housing Survey (Department for Communities and Local Government (2013)).

¹¹ In order to determine what properties to use to calculate a reference price, Axtell et al. (2014) developed the concept of ‘similar houses’, which are those sold recently whose quality is close to the sold house. This approach is therefore dependent on the estimation of house quality. Our study also made use of the house quality assumption, and therefore, the approach by Axtell et al. (2014) could be used.

a negative correlation with the number of days houses spend on the market. Fitting this equation to WhenFresh/Zoopla data also gave good agreement, for different parameter values (given in the Appendix).

If a house remains on the market from the previous time-step, its price is reduced with a 6% probability. Records of price reductions on WhenFresh/Zoopla showed a very good fit to a Gaussian in the log domain, so the reduction in offer price is modelled according to

$$p_t = \begin{cases} p_{t-1} (1 - \exp(\varepsilon)), & \text{with probability } 0.06 \\ p_{t-1}, & \text{with probability } 0.94, \end{cases} \quad (8)$$

where ε is a draw from a Gaussian distribution. If the price drops below the amount needed to pay the mortgage on the house, it is withdrawn from the market.

3.4.4 Decisions as a Buy-To-Let investor

Decision to buy a rental property. The decision is based on expected rental yield and expected capital gain. This set-up approximates different motives that investors may have, such as a steady stream of income or capital gains. In order to study the effect of the two different motives for investing, we have postulated two types of investors whose share we can experiment with. The first (second) type of investors put a weight δ on capital gain of 50 (90) percent, and a weight $(1 - \delta)$ on rental yield of 50 (10). The first type of investors can be understood as ‘fundamentalists’, since they value both streams of revenue equally, and the second type represent ‘trend-followers’, since they value mostly the capital gain and hence, whenever the market conditions are such that it is profitable to buy or sell properties, they are likely to all behave in the same way and hence follow the trend (see Section 6 for the experiments showing this behaviour).

BTL investors decide to add houses to their current portfolio based on the ‘expected yield’

on property investments, which is defined as

$$\begin{aligned}\Omega &= \text{leverage} \cdot (\delta \cdot \text{capital gain} + (1 - \delta) \cdot \text{rental yield}) - \frac{m}{d} \\ &= \frac{p}{d}(\delta(g + \kappa) + (1 - \delta)\bar{r}) - \frac{m}{d},\end{aligned}\tag{9}$$

where p/d is the leverage (purchase price p divided by downpayment d), δ is the weight on capital yield, g is the estimation of monthly house price growth defined in Equation (4), κ is the long-term, average gross yield, $(1 - \delta)$ is the weight on rental yield, \bar{r} is the current average gross yield in the market, and m/d is the mortgage rate (the monthly mortgage payment m divided by the downpayment d).

The probability of deciding to buy a house is then given by the logistic function

$$P_{buy,BTL} = \sigma(\beta\Omega)^{\frac{1}{12}},\tag{10}$$

where β is the shape parameter of the logistic function representing intensity of choice and the factor of $1/12$ is included because the decision is made every month. If the investor decided to buy a house, it will place a bid in the market for an amount equal to its maximum available mortgage,¹² and, if the bid is successful, it will choose a downpayment or possibly pay for the house outright (see the Appendix for details).

Buy-to-let rental offers. A BTL investor will put a house on the rental market whenever a rental contract ends, or when a new BTL house is bought or inherited that doesn't already have a tenant. BTL investors offer houses on the rental market at a monthly rental, given by

$$\ln q = \alpha + \ln(\bar{r}) - \beta \ln(\zeta(1 + \bar{f})) + \varepsilon.\tag{11}$$

This equation is of exactly same form as used for the sale price in equation (7), where the average

¹² Investors' bids are capped by the average market price of a house of the maximum quality; this prevents artificially large bids.

price \bar{p} is replaced by average rent \bar{r} , and \bar{f} is the number of days in the rental market, rather than in the ownership market.¹³ If a house on the rental market does not get filled, the price is reduced each month by a constant fraction of 5%. The length of a rental agreement is chosen randomly from 12 to 24 months with uniform probability, as suggested by data from the Association of Residential Letting Agents (ARLA).

Decision to sell a rental property. BTL investors will consider selling their investment properties at the end of each tenancy agreement, and will re-consider each month until another tenant moves in. The decision to sell is based on the ‘effective yield’ on the house, which is defined as

$$\begin{aligned}\Psi &= \text{leverage} \cdot (\delta \cdot \text{capital gain} + (1 - \delta) \cdot \text{rental yield}) - \frac{m}{u} \\ &= \frac{p}{u}(\delta(g + \kappa) + (1 - \delta)\bar{r}) - \frac{m}{u}\end{aligned}\tag{12}$$

which is the same as Equation (9), except that instead of the downpayment d , the house equity u is used, which is defined as the current market price of the house minus the mortgage principal, and is bound to be greater than zero. The probability of deciding to sell the house is then given by

$$P_{\text{sell},BTL} = 1 - \sigma(\beta\Psi)^{\frac{1}{12}}\tag{13}$$

which is equivalent to Equation (10) with opposite sign.

If an investor decides to sell, the house will be taken off the rental market and put on the sale market at the price given in Equation (7).

3.5 Housing markets

Market clearing is implemented as a double auction process, which proceeds in a number of rounds, where each round consists of two phases. In this section, *bids* refer to buyers seeking to buy homes

¹³ In addition, rent is bound by a minimum set at 4.8% of the house price, in order to prevent artificially low rents.

for up to a maximum price p_b , and *offers* (Q, p_s) refer to sellers seeking to sell their homes of quality Q for a minimum price p_s . Both buyers and sellers could be BTL investors or non-investor households.

In the first phase, all bids are matched to the offer with the highest quality Q that they can afford (i.e. having $p_s \leq p_b$). As a result, some offers will be matched to one or multiple bids while others might not be matched to any, and similarly, bids with too low a desired price p_b will not be matched to any offers.

In the second phase, the procedure iterates through the offers that were matched to any bids. When an offer was matched to a single bidder, that bidder is selected as the winner and the transaction is finalised at the price demanded by the seller. However, when a given offered house is matched with more than one bidder, the seller gets to ‘bid up’ the price by a small amount (see the Appendix A.4 for details) and out of the bids which can still afford the house, one is randomly chosen as the successful buyer.¹⁴

All unmatched bids and offers are returned to the pool, and the process continues in more matching rounds until (i) there are no bids left, (ii) there are no offers left, or (iii) the maximum number of rounds, which is set as a fraction of the population size, is reached. At the end of the clearing process, all remaining unmatched bids are deleted (the households who made them will have to bid again in the following month if they so decide) and all unsold offers remain, although the households who made them might decide in the following month to update their offer price, remove the offers from the market, or in the case of BTL investors, place them in the rental market instead.

Rental clearing proceeds in the same way as house-sales-clearing, with the only difference that rent is used instead of house price.

¹⁴ Note that we must choose the winner at random, and not simply the bidder with the highest bid, since the bidding price p_b (the maximum price a buyer is willing to afford) is secret and not known to the seller.

3.6 The bank

There is a single bank in the model which represents the mortgage lending sector in the aggregate. Potential owner-occupiers and BTL investors apply for mortgages by declaring to the bank their desired expenditure and down-payment. The approval processes for residential and BTL mortgages are explained below. Apart from providing mortgages, the bank offers deposit accounts as well and the monthly interest rate is 0.2%.

3.6.1 Mortgage approvals

Loans to owner occupiers. The bank offers 25 year fixed interest repayment mortgages to owner-occupiers. The bank will approve a mortgage to a (potential) owner-occupier as long as it conforms to loan-to-value (LTV), loan-to-income (LTI) and affordability constraints. The affordability constraint ensures that a household has enough total income to pay all its mortgages.¹⁵ Subject to meeting those criteria, all demand is met in any period. The maximum principal loan amount, then, is calculated as

$$q = \min \left(\frac{b\chi}{1 - \chi}, y\psi, y_d \nu \frac{1 - (1 + i_R)^{-N}}{i_R} \right). \quad (14)$$

The constraints are described in the following table.

¹⁵ In the current version of the model, this can be achieved by setting the interest rate at which affordability is assessed to the mortgage interest rate at origination. As all residential mortgage contracts are fixed interest repayment mortgages, this ensures that households can meet their mortgage payments during the course of the mortgage contract, holding income constant. In contrast, in a model with flexible interest rates, an affordability test could assume a “stressed” interest rate above the mortgage rate at origination. This would ensure that households could still afford the mortgage even if interest rates rose (up to a certain degree) during the course of the mortgage contract. In the current version of our model with exogenous monetary policy, the assumption of fixed rates is not as stringent as it may appear, though the effects of exogenous shocks to monetary policy on the housing market may be understated in this case. Extending the model to introduce floating interest rates is an avenue for further research worth exploring.

Constraint	Description
$\frac{b\chi}{1-\chi}$	LTV constraint. b is the household's bank balance (assuming all bank balance is used as downpayment), χ is the maximum loan to value ratio.
$y\psi$	LTI constraint. y is household gross income and ψ is the maximum loan to income ratio
$y_d\nu\frac{1-(1+i_R)^{-N}}{i_R}$	Affordability test given a monthly payment equal to the share ν of the household's disposable income available for mortgage payments. y_d is a household's disposable income, i_R is the fixed monthly interest rate ['R' stands for residential as opposed to BTL] and N is the number of monthly payments to pay off the mortgage.

Loans to buy-to-let investors. The bank offers 25 year fixed rate interest-only mortgages to BTL investors. The bank will approve a mortgage to a (potential) BTL investor as long as it conforms to LTV and interest cover ratio (ICR) constraints. The ICR limit imposes the constraint that

$$\Lambda < \frac{b}{1 - \frac{\bar{y}_r}{\xi i_{BTL}}}, \quad (15)$$

where \bar{y}_r is the moving exponential average of the gross annual rental yield on new tenancy agreements (i.e. gross annual rental income over house price), ξ is the ICR constraint applied by the bank and i_{BTL} is a mortgage interest rate. The ICR provides a buffer to ensure that borrowers can continue to afford mortgage payments if costs increase or rental income declines.

3.6.2 Interest rates

Mortgage interest rates are calculated as the bank rate plus the spread. Since the bank rate is exogenous, all the endogenous variation in interest rates comes from the spread. Mortgage interest rate spread, i^{spread} , is calculated each month according to

$$i_{t+1}^{spread} = i_t^{spread} + \alpha(M_t - T), \quad (16)$$

where α is a constant, M_t is the total supply of mortgage lending in month t and T is an exogenous constant target monthly supply. The rationale behind this equation is that if total lending by the bank exceeds its (exogenous) target supply, the bank will increase the mortgage interest rate to discourage further borrowing. Target supply is calibrated to match the UK data on aggregate credit.

This is a simplified way of modelling credit supply in a model without a profit-maximising banking sector. Introducing a more realistic (agent-based) banking sector is one of the main extensions of the model that has been left for further research at this stage.

3.7 The Central Bank

The central bank can set LTV, LTI, ICR and affordability policies. Policies can be of three different types:

1. Strict limits, e.g. a hard LTV limit of 90% for all households (though the limit may differ between types of agents, such as first-time buyers or owner-occupiers);
2. ‘Soft’ limits, e.g. an LTI cap of 3.5 on new mortgage lending, but allowing for 15% of new mortgages above this limit;
3. State-contingent policies, e.g. an LTV limit of 85% if credit growth over a certain time is above a certain threshold; otherwise no limit.

Monetary policy is exogenous in the model.

4 Calibration

There are approximately 27 million households in the UK. Due to computational reasons, the model is simulated at a reduced scale therefore contains 10,000 households, where one household in the model represents 2,700 real households.

The calibration of the model proceeds in two steps: a micro-calibration that fine-tunes household's behaviours or characteristics directly against suitable data, and a macro-calibration that ensures consistency with relevant economic aggregates (such as the FPC's core indicators). All parameters and their calibration are set out in detail in the Appendix B.

4.1 Micro-calibration

Whenever possible, individual components of the model have been calibrated directly with relevant data. A large set of micro data has been used for this purpose, mostly from household surveys and housing market data sources, e.g. PSD, Land Registry transactions data, and WhenFresh/Zoopla data on rental listings. It is important to notice that in the calibration of individual model components and equations, the target is to match *distributions* of data rather than aggregates or averages. This is one of the main strengths of an agent-based model. Our calibration therefore makes use of the whole distribution of the data, not just some of its moments. We also calibrate the model based on joint distributions of two or more variables wherever possible; that is, the distribution of age, income, mortgage amounts, property ownership, yields on investment properties etc. of the agents in our model matches the marginal and sometimes multivariate distribution of these variables in the UK population. Some examples of model equations that have been micro-calibrated include:

- the consumption equation has been calibrated in order to match the income and wealth distributions, obtained from financial household surveys;
- the life-cycle equations have been micro-calibrated directly against relevant statistical distributions: the mortality equation and birth rate were calibrated against ONS data, and the

income relationship was matched against the Wealth and Assets Survey.

4.2 Macro-calibration

Some equations or components of the model cannot be micro-calibrated, since it was not possible to obtain suitable data sets against which to match the parameters. In these cases, the approach followed is to ‘macro-calibrate’ the parameters in order to obtain a desired output.

First, a suitable objective function or performance metric must be defined. A common approach followed in other ABMs (for example, Geanakoplos et al. (2012)) is to train the model on a time series covering a few years and attempt to predict the following few years. This approach was not followed in this study because our model does not seek to represent any specific point in time in the UK housing market, but instead tries to model the baseline or long-run dynamic of the market.

Instead, the performance of our model has been defined as the agreement with the average values of the FPC’s core indicators. The set of nine core indicators defined by the FPC are considered to provide a comprehensive picture of the state of the housing market at a given time, and by matching their averages we ensure that our model represents a realistic state of the housing market.

This method is less desirable than the micro-calibration, since it is approximate and requires some intuition from the modeller, so it has only been used when micro-calibration was not possible. Examples of the model equations that have been calibrated in this way include:

- the selling rule has been set to match the average moving frequency for owner-occupiers, while preventing unrealistic build-up of housing stock;
- the renting vs. buying behavioural rule has been set to achieve the right distribution of renters versus owner-occupiers.

5 Validation

One of the key features of the model is its ability to generate real-world-like house price cycles which come naturally from the interactions of buyers and sellers in the market. Due to the stochastic nature of the model, we rely on Monte Carlo simulations for the policy experiments. That is, we simulate the model multiple times, which gives us a distribution for each variable.


To illustrate the endogenous cycles generated by the model, we show the house price index in a sample simulation run of 45 years (after initialising the model for approx. 200 years) in Figure 2a. The main pro-cyclical factor is the expectations component, encapsulated in the house price growth expectation defined in Equation (4), that leads to BTL investors buying in rising prices and selling in falling prices. We can show this is the case by running a simulation in which we set the house price expectation g to zero. As can be seen in Figure 2b, shutting down the expectations channel almost completely removes the cyclical behaviour.¹⁶

More broadly, the main empirical facts the model is trying to replicate are the FPC's housing core indicators, listed in Figure 3. The figure shows that the model simulations are in mostly line with the historical data. It is worth noting that:


- real credit growth and real house price growth are zero in the model as they are long-term averages and there is no economic or population growth in the model;
- the owner-occupier debt-to-income ratio and house-price-to-income ratio are higher than what is observed in reality because of differences in the way income is calculated.

We perform numerous checks to assess the validity of the model. Due to space constraints, we only show some of these checks here. They should be seen as qualitative as the model is not matched to a specific point in time, but rather a long-run representation based on current economic

¹⁶ There are still small fluctuations in the model due to market frictions (for example, when a BTL investor cannot sell a property because the rental contract has not ended yet, or a seller cannot find a buyer and has to lower the price of the property in the following month) as well as fluctuations due to the noise terms in the equations and the randomness in the probabilistic behavioural rules.



(a) Benchmark case, displaying boom and bust cycles



(b) Experiment with house price growth expectation g set to zero, no longer showing any boom and bust cycles

Figure 2: Evolution of house price index for two different simulation runs of 45 years (after initialising the model for approx. 200 years). At the top, the benchmark case using the house price growth expectation of Equation (4) creates endogenous boom and bust cycles. At the bottom, we artificially shut down the expectations channel by setting the expected house price growth to zero, and as a result, the cyclical behaviour disappears.

circumstances and policies (such as the interest rate or macroprudential policies). As such, we evaluate whether the model broadly matches key aspects observed in the data and relationships between variables based on economic theory.

In a first test, we show the correlation between house price growth and credit growth in Figure 4. The model does not directly model the relationship between these two variables - instead, the relation arises endogenously due to the behaviour and interactions of the agents in the model. The

Core Indicator	FPC's Policy Statement on housing tools			Housing ABM simulations since 1987
	Minimum since 1987	Average 1987-2006	Maximum since 1987	
OO mortgage LTV ratio (mean above the median)	81.6%	90.6%	90.8%	82.5%
OO mortgage LTI ratio (mean above the median)	3.6	3.8	4.1	3.9
BTL LTV ratio (mean)	70.9%	n.a.	78.6%	75.3%
Household credit growth	-0.1%	10.3%	19.6%	0.0%
OO debt to income ratio	72.8%	86.1%	105.4%	127.6%
Mortgage approvals	26,658	97,940	135,579	84,516
Advances to homemovers	14,300	48,985	93,500	52,499
Advances to FTB	8,500	39,179	55,800	21,770
Advances to BTL purchasers	3,603	9,903	16,230	10,247
House price growth	-5.6%	1.8%	7.0%	0.0%
House price to income ratio	2.3	3.2	5.0	5.3
Rental yield	4.8%	5.8%	7.6%	4.9%
Spreads (basis points)	34	81	361	251

Figure 3: FPC's Housing Core Indicators

Note: For definitions of the variables, see the FPC's Policy Statement on its powers over housing tools: <http://www.bankofengland.co.uk/financialstability/Documents/fpc/policystatement010715.pdf>, p.33. Color coding provides a rough indication of whether our simulation results are in line with empirical data.

figure shows that episodes of higher credit growth are associated with higher house price growth: a one percent increase in credit growth translates into a roughly 0.13 percent increase in house price growth. This is what we would expect: expansions in credit and housing booms tend to go together, though there is a considerable amount of uncertainty around this relationship. Though not directly comparable, our estimate is close to the study by Favara and Imbs (2015), who found that, in the US, a one percent change in credit increases the growth rate of house prices by 0.2 percent.

We also perform experiments unrelated to the policy variables of interest to see whether the variables in the model behave as expected. Figure 5 shows the effects of several such experiments on the housing core indicators.

The changes have the expected sign. For example, a general exogenous increase in households income by 10% would reduce the need for mortgage advances (decreases of 5 to 9%, depending on the type of household) and reduce overall debt burdens (overall fall of 10%). Similarly, an increase in the banking sectors credit target by about 30% increases advances by between 0 and 43% (depending on the type of household), mortgage approvals by 30%, and debt burdens by 8%. Interest rate spreads decline by 4% as there is more credit supply in the market. An increase in housing supply by 7% causes higher advances to BTL investors (22%) and first-time buyers (1%),

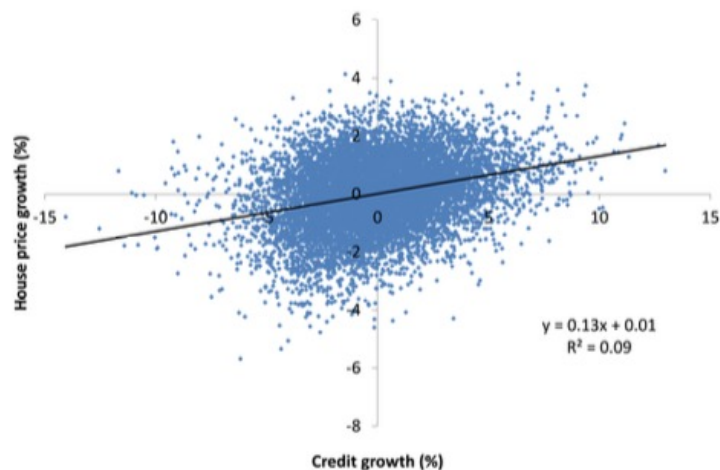


Figure 4: Endogenous relationship between house price and credit growth

Core indicator \ Experiment	Increase in income by 10%	Increase in credit target by 29%	Increase in housing supply by 7%
BTL LTV	1.00	0.99	0.98
Owner-occupier LTI	1.01	0.98	1.00
Advances to BTL	0.91	1.33	1.22
Advances to first-time buyers	0.95	1.00	1.01
Advances to movers	0.92	1.43	1.00
Debt-to-income ratio	0.90	1.08	1.03
Housing transactions	1.02	1.33	1.01
Interest rate spread	1.00	0.96	1.00
Mortgage approvals	0.92	1.30	1.03
Owner-occupier debt-to-income	0.90	1.06	1.00
Price-to-income ratio	1.00	0.97	0.97
Rental yield	0.99	0.97	0.97

Figure 5: Effect of various experiments on housing core indicators relative to benchmark (benchmark = 1)

though there is no change for home movers. This makes sense as the number of properties going to first-time buyers and home movers is limited, given a constant population size, while BTL investors can add to their portfolio. The increased supply of BTL properties puts downward pressure on rental yields, which fall by 3%.

Another validation test is to compare empirical distributions to those produced by the model. Figures 6 and 7 show the distribution of mortgages by loan-to-income and loan-to-value bands in the PSD and the model.

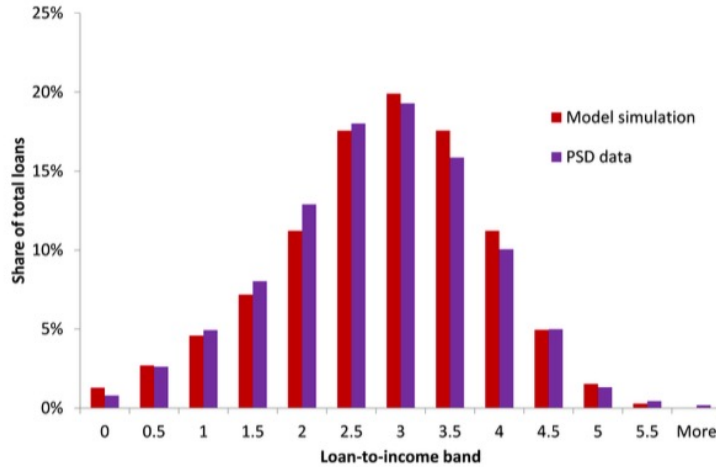


Figure 6: LTI distribution

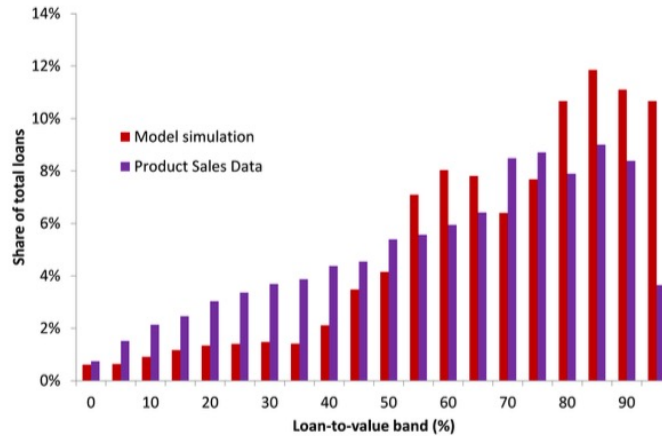


Figure 7: LTV distribution

6 Experiments

The calibrated model is used to perform a series of comparative statics exercises to investigate the impact of the size of the rental/ BTL sector and different types of BTL investors on housing booms and busts (Section 6.1). It is also used to qualitatively assess the effect of macro-prudential policies, such as an LTI limit (Section 6.2).

6.1 Buy-to-let market

This section presents a number of experiments performed to understand the importance of (i) the size of the BTL market and (ii) different types of BTL investors regarding house price volatility and the amplitude of house price cycles.

6.1.1 Size of the buy-to-let market

To recap, BTL investors' decision to buy or sell houses is based on the "expected yield" on property investments, which is a function of expected annual house price appreciation, the expected gross annual rental yield, the annual mortgage payment and the minimum down-payment. The higher (lower) the expected yield, the more likely it is that BTL investors buy (sell) a rental property. There are two types of investors, in line with the baseline calibration of the model: The first (second) type of investors put 50 (90) percent weight on capital appreciation and 50 (10) percent weight on rental yield. The share of first and second types of BTL investors can be changed, and is set to equal in the benchmark case.

To investigate the effect of the number of BTL investors, two scenarios are considered: (i) a benchmark scenario of around 4 percent of households being BTL investors¹⁷, and (ii) an alternative scenario of around 16 percent of households being BTL investors. These two scenarios are attained by setting the BTL gene parameter exogenously. In the benchmark scenario, the calibration results in BTL investors holding 25% of the housing stock. In the alternative scenario, investors hold 48% of the housing stock.

Figure 8 displays the distribution of average annual house price growth rates during boom and bust episodes for these two scenarios. We first smooth the simulated house price growth series to eliminate small fluctuations. We then identify boom episodes that last from a market trough to the market peak, and bust episodes that last from a market peak to a market trough. We finally

¹⁷ Combining the number of private rental houses (ONS, 2011) with the distribution of number of houses owned by BTL investors (ARLA, 2014) suggests 4 percent of population are BTL investors.

calculate the average growth rate during each episode. This measure is better suited to show the magnitude of the house price cycles, as opposed to alternative measures, for example the standard deviation. The left panel shows the benchmark case, and the right panel shows the case where the share of BTL investors is larger, resulting in the frequency of sharper house price movements being higher. An alternative measure, the standard deviation of quarterly house price growth, shows a similar result as it increases from 1.2 to 2.3 percent.

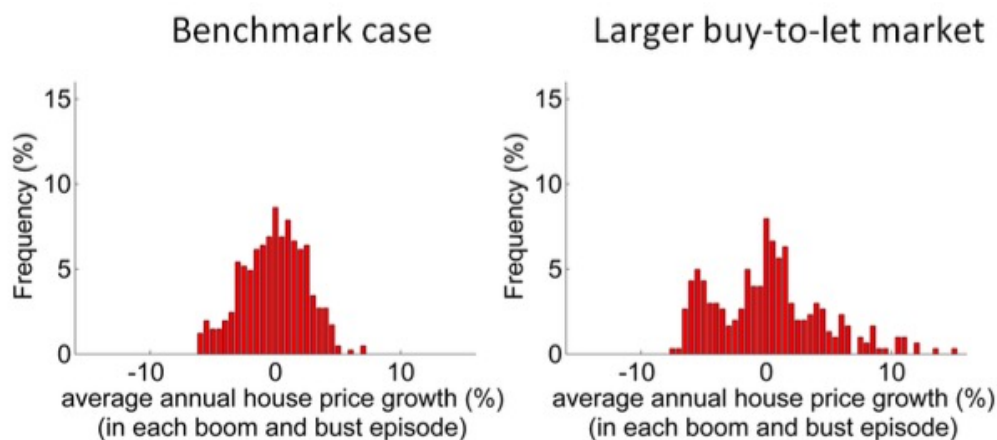


Figure 8: The distribution of average annual house price growth rates in booms and busts

This result is driven by the presence of more buyers and sellers in the market whose decisions are more responsive to market conditions. This amplifies the house price cycles. During booms the increased demand from BTL investors drives up house prices, and during busts the increased supply from BTL investors drives down prices further. The key assumption behind this amplification channel is the way house price expectations are formed: households extrapolate last year's house price growth rate when determining the expected yield on property investments. When we turn off the expectations channel in the model by setting the expected house price appreciation g to zero (see Equation (4)), the size of the BTL market does not have a significant impact on house price volatility and on the amplitude of the house price cycles.

The increased demand from BTL investors also drives up the level of house prices by around 59 percentage points. This higher price reduces the number of owner-occupiers that can afford to

buy a house, reducing their market share by about 25 percent. They turn into tenants, renting from the BTL investors that bought the houses. These results should be seen as a comparative statics exercise - they do not answer the question what the optimal size of the BTL sector should be in the housing market. This would require a measure of welfare that is currently not present in the model.

6.1.2 Different types of buy-to-let investors

Furthermore, to investigate the role of different types of investors on house price volatility and housing cycles, two sets of experiments are conducted: (i) all BTL investors are calibrated to only care about rental yield, and (ii) all BTL investors are calibrated to only care about capital gain. The difference in these two experiments is not very significant when the BTL sector is around its current size of 4 percent (results not shown).

However, the impact of the composition of BTL investors may depend on the size of the market. Therefore, the same experiment has been conducted when the size of the BTL sector is increased to 16 percent. Figure 9 displays the distributions of average annual house price growth during boom and bust episodes when BTL investors are calibrated to only care about rental yield (left panel), and when BTL investors are calibrated to only care about capital appreciation (right panel). With the larger BTL market, we find that the market volatility, quantified by the frequency of sharp house price movements and by the standard deviation of annual house price growth, are much higher when investors care about capital appreciation than when they care about rent, as we would expect.¹⁸

The destabilising effect of BTL investors' sensitivity to capital gain is not surprising: in contrast to owner-occupiers, BTL investors follow trends and have the ability to own (and potentially sell) multiple properties. Therefore, as the number of investors increases, the market will increasingly be dominated by their behaviour. This augments the size of booms and busts. The three-peaked house-price-growth histogram on the right shows that there is a critical rate of growth

¹⁸ When BTL investors are calibrated to only care about rental yield (capital gain), the standard deviation of house price growth is 1.7 (2.6).

at which investors begin to jump on the bandwagon, in the hopes of a good return from capital gain, and start a large boom. Equally, there is a critical rate of depreciation at which investors begin to abandon ship and start a large bust.

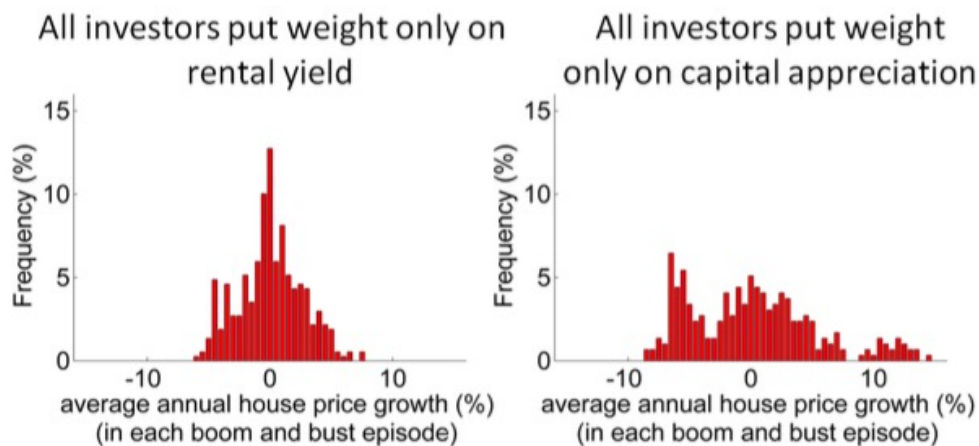


Figure 9: The distribution of average annual house price growth rates in booms and busts: the role of different types of BTL investors with a larger BTL sector

However, we note, by comparing the left panel of Figure 9 to the original benchmark case in the left panel of Figure 8, that the BTL investors motivated by rent are actually introducing a *stabilising* effect on house prices. This effect on the ownership market takes place via the *rental* market, in the following way. When HPI is increasing, investors tend to buy more properties, which increases the supply on the rental market, putting a negative pressure on rental prices, and hence reducing the likelihood that renters will decide to become owner occupiers (and also making it more likely that owners will switch to renting). The result is a decrease in demand in the ownership market, which is a negative feedback (stabilising) effect.

Therefore, we conclude that the BTL sector is exerting two competing effects on the system: (i) a pro-cyclical feedback effect linked to capital appreciation (buying in rising prices and selling in falling prices), and (ii) a stabilising feedback effect linked to renting income which increases the size of the rental market and reduces the size of the ownership market in rising prices. The relative strengths of these two competing effects depend on the size of the BTL sector. At 4% BTL, the two effects are comparable, whereas at 16% BTL, the pro-cyclical feedback is much larger because the

proportion of renters in the whole system is higher.

6.2 Macprudential Policies

The model is also well suited to assess the effects of macroprudential policies. Such measures are designed to make the financial system as a whole safer. But the immediate targets of these policies are individual entities, such as households or banks. It is straightforward to implement such policies in an ABM at the level of the individual and observe the aggregate effects they have on the system. Other types of models often miss the effects of such policies. For example, it is difficult in representative agent models to assess the effects of policies that have a larger impact on some individuals than on others. For instance, a policy that drives some but not all individuals out of the market requires a degree of heterogeneity among agents that is not straightforward to achieve in representative agent models.




Figure 10: Distribution of mortgages by LTI band with and without LTI experiment

To highlight these benefits of the housing ABM, we conduct a policy experiment where we impose a regulatory restriction on the agents' ability to access a mortgage. In this case, we implement a loan-to-income limit of 3.5, but allow the lender to issue 15 percent of mortgages with an

LTI above this limit.¹⁹

The Figure 10 shows the share of lending for each bucket of width 0.5 of the LTI distribution. Most mortgages have an LTI ratio of 2 to 4, reflecting the fact that most households need to take out a substantial mortgage in order to buy their first home or move to another one. In our simulations, we do not observe mortgages with very high LTI ratios as they are subject to other criteria as well, such as lender-imposed loan-to-value limits and debt-service ratios.

In the case of an LTI limit of 3.5, the policy substantially reduces the incidence of high LTI mortgages. There is clustering right below the soft limit of 3.5 as many aspiring first-time buyers and home movers are pushed below the limit.

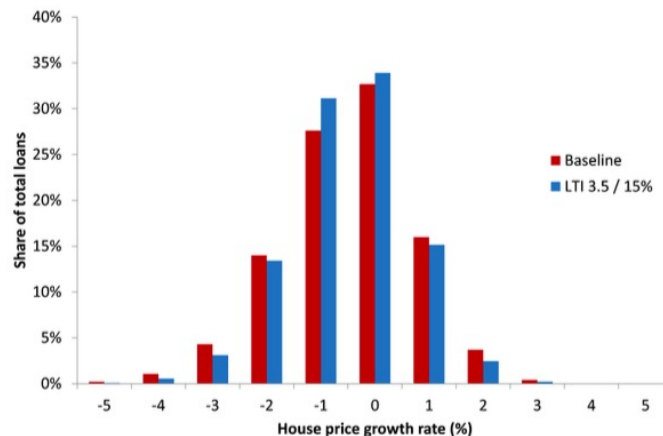


Figure 11: Distribution of house price growth (in percent) with and without LTI experiment

We can also analyse the effect on every other variable in the model as well, especially the housing core indicators. For example, Figure 11 displays the baseline and policy scenario in a histogram of the growth of house prices. The most striking aspect is that the distribution under the policy experiment has a lower standard deviation than the baseline case: it decreases from 1.21 to 1.09. Hence, the LTI limit has a dampening effect on the boom-bust cycle of house prices.

¹⁹ In this exercise loan-to-income limit of 3.5 is chosen for illustration only. The Bank of England's loan-to-income flow limit that came into force in October 2014 set at 15 percent of new mortgages at an LTI multiple at, or greater than, 4.5. See <http://www.bankofengland.co.uk/prd/Documents/publications/ps/2014/ps914.pdf>.

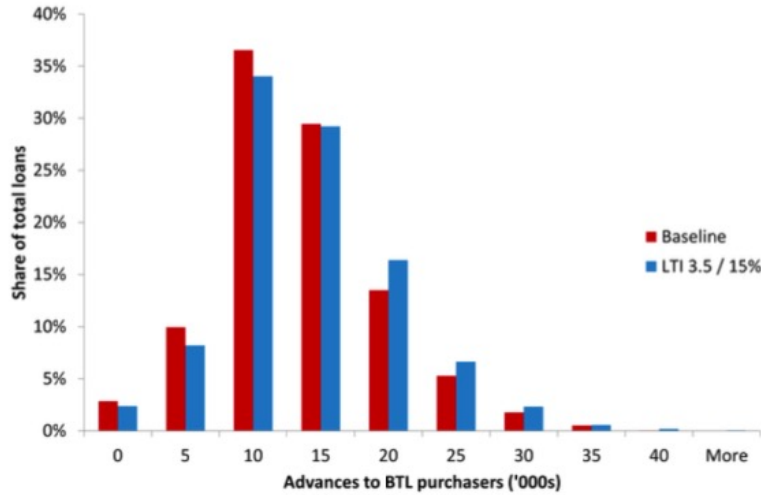


Figure 12: Distribution of advances to BTL purchasers with and without LTI experiment

Moreover, we can assess the effects of the policy on different market segments. For example, Figure 12 shows that the BTL sector grows in response to the LTI limit that is imposed on owner-occupiers, but not investors. Advances to BTL purchasers of between 20,000 and 30,000 increase especially strongly. On average, advances to BTL purchasers increase from 10,247 to 11,037, or by about 7.7%.

We can also easily implement stage-contingent policy experiments. This could reflect the fact that macroprudential policies may be designed to depend on the economic cycle, rather than being structural measures. For example, we can impose the previous LTI policy, but only when credit growth is above a certain threshold. In this case, the mass in each bucket of the histograms lies between the baseline case and the LTI policy that applies independently of credit growth (not shown).

These results can help policymakers get a better understanding of the effects of their policies on the market as a whole. In particular, it allows them to look at different market segments and see changes to the whole distribution of their variables of interest, which is a natural result of the heterogeneity of the agents in the model.

7 Conclusion

This paper has demonstrated the potential usefulness of agent-based modelling for policy-making purposes in the context of the UK housing market. We have highlighted the strengths of ABMs by exploiting the heterogeneity of the modelled agents for an LTI limit of 3.5 with a 15% allowance of mortgages above the limit, showing the existence of realistic non-linear dynamics, and using a large number of micro-datasets to get at a fine-tuned calibration of households' behavioural equations. The model can shed light on the dynamics in the BTL sector caused by different behaviours of agents in the market.

Our model can be expanded to assess the effects of even more policy options. Further research into incorporating an agent-based banking sector would allow us to model the interaction between mortgage lenders and borrowers in greater detail, so we could introduce changes in capital requirements as another instrument for policy-makers to affect the mortgage or housing markets. Additionally, this would enable us to model the mortgage product space in a more realistic way, including a variety of fixed- and floating-rate mortgage products with different terms and fee structures. Last but not least, introducing a spatial dimension in the model could allow us to analyze regional differences in the housing market, such as London compared to other parts of England.

ABMs may still be at an early stage in terms of their use for economic modelling, but our results presented in this paper show the potential of this approach as an additional tool for policy-making.

References

- Axtell, R., D. Farmer, J. Geanakoplos, P. Howitt, E. Carrella, B. Conlee, M. Hendry, P. Kalikman, D. Masad, N. Palmer, et al. (2014). An agent-based model of the housing market bubble in metropolitan washington, dc.
- Bernanke, B. S., M. Gertler, and S. Gilchrist (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of macroeconomics 1*, 1341–1393.
- DCLG - Department for Communities and Local Government (2014). English Housing Survey, 2012-2013: Household Data [computer file]. Technical report, Colchester, Essex: UK Data Archive.
- Department for Communities and Local Government (2013). English Housing Survey, 2008-2013: Housing Stock Data [computer file]. Technical report, Colchester, Essex: UK Data Archive.
- Devisch, O. T. J., H. J. P. Timmermans, T. A. Arentze, and A. W. J. Borgers (2009, August). An agent-based model of residential choice dynamics in nonstationary housing markets. *Environment and Planning A 41*(8), 1997–2013.
- Diappi, L., B. P. (2008). Smith’s rent gap theory and local real estate dynamics: A multi-agent model. *Computers, Environment and Urban Systems 32*(1), 6–18.
- Erlingsson, E., A. Teglio, S. Cincotti, H. Stefansson, J. Sturlusson, and M. Raberto (2014). Housing market bubbles and business cycles in an agent-based credit economy. *Open-Assessment E-Journal 8*, 1–42.
- Farmer, J. and J. Foley (2009). The economy needs agent-based modelling. *Nature 460*, 685–686.
- Favara, G. and J. Imbs (2015). Credit supply and the price of housing. *The American Economic Review 105*(3), 958–992.

- Funke, M. and M. Paetz (2012). A dsge-based assessment of nonlinear loan-to-value policies: Evidence from hong kong.
- Ge, J. (2014). *Who Creates Housing Bubbles? An Agent-Based Study*, pp. 143–150. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Geanakoplos, J., R. Axtell, J. D. Farmer, P. Howitt, B. Conlee, J. Goldstein, M. Hendrey, N. M. Palmer, and C.-Y. Yang (2012, May). Getting at Systemic Risk via an Agent-Based Model of the Housing Market. *American Economic Review* 102(3), 53–58.
- Gelain, P. and N. Bank (2011). Macro-prudential policies in a dsge model with financial frictions. In *7th Dynare Conference, Atlanta, USA*.
- Gelain, P., K. J. Lansing, and C. Mendicino (2012). House prices, credit growth, and excess volatility: Implications for monetary and macroprudential policy.
- Gilbert, N., J. C. Hawksorth, and P. A. Swinney (2009). An agent-based model of the english housing market. In *AAAI Spring Symposium: Technosocial Predictive Analytics*, pp. 30–35.
- Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *The American economic review* 95(3), 739–764.
- Iacoviello, M. and S. Neri (2010). Housing market spillovers: evidence from an estimated dsge model. *American Economic Journal: Macroeconomics* 2(2), 125–164.
- Jordan, R., M. Birkin, and A. Evans (2010). Agent-based simulation modelling of housing choice and urban regeneration policy. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pp. 152–166. Springer.
- Kannan, P., P. Rabanal, A. Scott, et al. (2009). *Monetary and macroprudential policy rules in a model with house price booms*. International Monetary Fund, Research Department.

- Kouwenberg, R. and R. C. Zwinkels (2015). Endogenous price bubbles in a multi-agent system of the housing market. *PloS one* 10(6), e0129070.
- Lambertini, L., C. Mendicino, and M. T. Punzi (2013). Leaning against boom–bust cycles in credit and housing prices. *Journal of Economic Dynamics and Control* 37(8), 1500–1522.
- Magliocca, N., E. Safirova, V. McConnell, and M. Walls (2011). An economic agent-based model of coupled housing and land markets (CHALMS). *Computers, Environment and Urban Systems* 35(3), 183–191.
- McMahon, M., A. Berea, and O. H. (2009). An Agent Based Model of the Housing Market. *Computers, Environment and Urban Systems*.
- Mendicino, C. and M. T. Punzi (2014). House prices, capital inflows and macroprudential policy. *Journal of Banking & Finance* 49, 337–355.
- Office for National Statistics (2014). Social Survey Division, Wealth and Assets Survey, Waves 1-3, 2006-2012 [computer file]. Technical report, Colchester, Essex: UK Data Archive.
- Office for National Statistics and Department for Environment, Food and Rural Affairs (2014). Living Costs and Food Survey, 2008-2013 [computer file]. Technical report, Colchester, Essex: UK Data Archive.
- Pangallo, M., J. P. Nadal, and A. Vignes (2016, June). Price formation on a housing market and spatial income segregation. *ArXiv e-prints*, arXiv 1606.00424.

A Appendix: detailed model descriptions

A.1 Simulation initialisation

In order to initialise a simulation with a realistic assignment of houses and mortgages to households, the model goes through a ‘spin-up’ period at the beginning of every simulation. The spin-up period begins with no households and no houses. While the number of households in the model is below the target population size, new households are added at an increased birth rate, while they age and die according to the usual mortality distribution described in Section 3.2.2. Once the target population has been reached, the birth rate is set to the correct value, which together with the mortality distribution guarantees a stable population.

During spin-up, new houses are put on the sale market by a ‘construction sector’ until the household to house ratio reaches the target value²⁰. New houses are put onto the house sale market at a price based on the ONS house price index data tables for 2013. If unsold, the price is reduced at a rate of 5% per month until a household buys them.

As a result, at the end of the spin-up period, the target population size, age distribution, and number of houses have been reached. We typically allow the model to spin up for 200 years. For the purposes of processing the simulation output, the spin-up part of the simulation is discarded when computing any statistics.

A.2 Downpayments

Whenever a household decides to buy a house and wins a bid in the market, it must first choose a downpayment. A relationship between downpayment and income percentile has been obtained from PSD. Examining the data revealed significant differences between the relationships for first-time buyers and other owner-occupiers.

For all owner-occupiers (including first-time buyers), the desired downpayment d depends on

²⁰ Household to privately-owned house ratio is 1.22 (Office for National Statistics (2014))

the income percentile of the household: households in the bottom 30% of income have a desired downpayment of 0, and for the rest of the households, downpayment is related to their income percentile according to the log-normal distribution

$$d = \max(d_{\text{minimum}}, \text{HPI} \cdot F^{-1}((\text{Income Percentile} - 0.3)/0.7)), \quad (17)$$

where d_{minimum} is the minimum downpayment set by the bank, HPI is the house-price index, F is a Log-Normal distribution with parameters (m, s) and F^{-1} is the inverse cumulative distribution. The distribution's parameters (m, s) are different for first-time buyers and other owner-occupiers, and the parameter values can be found in Section B of the Appendix.

For the buy-to-let investors' downpayment, we assumed a normal distribution of downpayment fractions, and set the parameters so that the LTV ratios are within the expected range

$$d = \max(d_{\text{minimum}}, p \cdot N(\mu, \sigma)). \quad (18)$$

A.3 Market price information

Households make use of market price information when making their decisions. For both the rental and sale markets, two types of information are available: the house price index and the market price of a house of a given quality.

The house price index for a given month is defined as the average transaction price divided by the average reference price over the set of all completed transactions for that month. The reference price of a house is the price of a house of that quality according to the ONS house price data tables 2013. This is an approximation to Nationwide's mix adjustment methodology. Since the houses in the model have only one property, namely quality, the hedonic regression can be reduced to a regression on $p = \chi p_r$ where p is the price of a house now, χ is the house price index and P_r is the reference price, which can be thought of as a measure of quality.

The market price given quality is calculated as a moving exponential average of completed

transactions involving houses of the given quality. Because the number of transactions per month may be quite small in the simulation (due to scaling down of the population) some quality bands may have very few transactions which leads to unrealistic distributions of price with quality. Analysis of house price distribution data over time shows that the shape of the distribution stays the same (almost log-normal). When the simulated population is large the model naturally reproduces this.

To deal with scaled-down populations, however, at the end of every month the market price given quality is transformed according to

$$p'_q = Dp_q + (1 - D)hp_r(q),$$

where D is a constant, p_q is the market price of quality q , h is the house price index for this month and $p_r(q)$ is the reference price of houses of that quality. This effectively relaxes the distribution of house prices to the shape (but not the level) of that in the 2013 data tables.

A.4 Sales market clearing

Clearing proceeds as follows: Home-buyers are matched to the best quality house they can afford and BTL investors are matched to the best expected-gross-rental-yield house they can afford. The expected-gross-rental-yield of a house of quality q , for-sale at a price p is defined as

$$E(y_q) = \frac{12\bar{r}_q E(o_q)}{p}, \tag{19}$$

where \bar{r}_q is the current market monthly rental price of houses of quality q and $E(o_q)$ is the expected occupancy of a rental property of quality q , (i.e. the expected fraction of time that rent will be collected on the property); this is based on an 18 month rental contract followed by a number of

days waiting for the next tenant, giving

$$E(o_q) = \frac{547}{547 + \overline{D}_q}, \quad (20)$$

where \overline{D}_q is the exponential moving average of the number of days that newly rented properties of quality q spent on the rental market.

When a given offered house is matched with more than one bidder, the price is ‘bid up’ by multiplying by 1.0075^k where k is chosen at random from a geometric distribution such that

$$p(k) = (1 - e^{-7b/30})^{k-1} e^{-7b/30}, \quad (21)$$

where b is the number of bids received in the timestep. The house is then offered to a randomly chosen bid that can still afford to buy. This approximates the outcome that would be achieved if the bids came in on random days in the simulated month; if a bid is followed by another bid within 7 days, the new bid ‘bids up’ the price by 0.75%, the first bid that is not bid up within 7 days gets the house.

Failed bids then get to bid again. This re-bidding carries on up to the smaller of $N/1000$ and $1 + n/5000000$ times, where N is the population and n is the total number of orders on the market.

B Appendix: parameters, initial conditions and data sources

Model component or Equation	Parameter values	Sources	Notes
Demographics			
Number of Households	10,000	Model input	
Birth rate	1.02%	English Housing Survey (DCLG - Department for Communities and Local Government (2014))	
Mortality		ONS Statistical Bulletin: Historic and Projected Mortality. Data from the Period and Cohort, Life Tables, 2012- based, UK, 1981-2062.	The pdf was multiplied by a constant factor so that the overall death rate is set equal to the birth rate, in order to ensure a constant population in the model
Income and financial wealth			
Income given age and income percentile		Living Costs and Food Survey (Office for National Statistics and Department for Environment, Food and Rural Affairs (2014))	
Minimum income	£5900	www.nidirect.gov.uk	Married couple's monthly lower earnings from income support
Essential consumption fraction	80%		Percentage of the minimum income spent by every household each month as "essential consumption"
Equation (1) - Desired bank balance	$\alpha=-32.00, \beta=4.07, \epsilon=N(0,0.1)$	Wealth and Asset Survey (Office for National Statistics (2014))	ϵ is constant for each household and represents a "propensity to save"
Equation (2) - Consumption fraction	C=0.5		Fraction of the available monthly budget the household uses for non-essential consumption
Return on financial wealth	0.20%		Interest rate for bank deposits
National Insurance	NI bands (NI rates) = £7755 (12%), £41450 (2%)	Government figures for 2013/14	
Income tax	Tax bands (tax rates) = £9440 (20%), £9440+32010 (40%), £9440+150000 (45%)	Government figures for 2013/14	
Behavioural rules			
Equation (3) - Expenditure	$\alpha=4.5, \beta=0.08, \epsilon=N(0,0.5)$		

Equation (4) - Price variation expectation	$\alpha=0.5$	Value suggested by John Muellbauer	α represents the estimation of the trend. $\alpha=1$ implies that the same trend is predicted.
Equation (5) - Buy or Rent decision	$\beta=1/3500; \tau=1.1/12$		β is the sensitivity of the rent/buy decision with respect to the cost difference. τ is the psychological cost of renting
Renting fraction	33%		fraction of income bid as rent
Equation (6) - Selling decision (home-owners)	Long term average= 11 years, $\alpha=4.0, \beta=5.0$	English Housing Survey (Department for Communities and Local Government (2013))	α is a penalty to excessive number of houses for sale, β is a penalty to excessively high interest rates
Equation (7) - Price demanded	$\alpha=0.04, \beta=0.011, \tau=1.0/31.0,$ $\epsilon=N(0,0.5)$		
Price decline if unsold	$P=5.5\%, \epsilon=N(1.603,0.617)$	WhenFresh/Zoopla	
BTL capital yield (trend followers)	$\delta=50$		
BTL capital yield (fundamentalists)	$\delta=90$		
Share of trend followers	[50%]	Model input	
Equation (10) - BTL purchase decision	$\beta=50.0$		Sensitivity or intensity of choice
Downpayment in cash	2	Calibrated against mortgage approval/housing transaction ratio, core indicators average 1987-2006	If the ratio of bank balance to house price is above this, payment for the house will be made in cash
Equation (17) - Downpayment for first-time buyers	$m=10.30, s=0.9093$	PSD	m is the scale parameter and s is the shape parameter of the log-normal distribution
Equation (17) - Downpayment for Owner-Occupiers	$m=11.155, s=0.7538$	as above	as above
Equation (11) - BTL Rent demanded	$\alpha=0, \beta=0, \zeta=1.0/31.0,$ $\epsilon=N(0,0.05)$		
Price reduction on rental market	5.0%		reduction in demanded rent each month until occupied
Equation (13) - BTL sale decision	$\beta=50.0$		Sensitivity or intensity of choice
Housing Markets			
Initial HPI	0.8		
HPI Log-Normal distribution	$m=195000, s=0.555$		scale and shape parameters

Profit margin for BTL investors (rent gross yield)	5%	Bank of England, unpublished analysis based on When-Fresh/Zoopla and Land registry matching
Length of rental agreement	Uniform, (12-24) months	ARLA - Members survey of the Private Rented Sector Q3 2013
Bank		
Initial bank base rate	0.50%	
Maximum LTV	0.9	Maximum LTV for Owner-Occupiers when unregulated
Maximum BTL LTV	0.8	Maximum LTV for BTL investors when unregulated
Maximum LTI	6	Maximum LTI for Owner-Occupiers when unregulated
Affordability Coefficient	0.5	Maximum proportion of income to be spent on mortgage
BTL interest	0.05	Interest rate for BTL when calculating ICR
Credit Supply Target	380	Target supply of credit per household per month
Initial mortgage interest rate	2%	
dDemand/dInterest	1.00E+11	Rate of change of credit demand with respect to interest rate (used to recalculate interest rate)