An Agent-Based Model of the English Housing Market

Nigel Gilbert
Centre for Research in Social Simulation
University of Surrey
Guildford, GU2 7XH, Surrey, UK
n.gilbert@surrey.ac.uk

John C. Hawksworth and Paul A. Swinney
PricewaterhouseCoopers LLP
1 Embankment Place
London WC2N 6RH, UK

Abstract

Markets for domestic housing in countries like England differ from other markets because the stock is fixed in the short term, buyers need to raise mortgages to finance their purchases and prices tend to be set through intermediaries ('estate agents' or realtors). We introduce an agent-based model in which some of the main features of the English market emerge from interactions between buyers, realtors and sellers and use this model to investigate shocks to the market.
Key words Housing market, Agent-based model, Realtors, House prices

The housing market

The current 'credit crunch' has highlighted the importance of the housing market in advanced economies such as England, where investments in housing are a major component of personal wealth. The conventional 'top down' approach is to estimate an econometric model that produces quantitative house price projections based on trends in incomes, interest rates and housing supply and demand (see PricewaterhouseCoopers (2007) for an example). However, econometric methods provide little understanding of the mechanics of the housing market.

In this paper, we use an agent-based model (ABM) to simulate the interaction of individual households who seek to buy and sell properties. House prices emerge from this market process. The potential advantages of using an ABM compared to econometric techniques include the ability:

- to model location, which is critical to house prices at a micro-level. Conventional econometric housing models have no spatial dimension, whereas this is a key feature of ABMs;

- to incorporate estate agents—called 'realtors' here to avoid confusion with the agents in the simulation—who play an important intermediary role in setting house prices; and

- to explain better the variation in house price to earnings ratios over time.

There have been surprisingly few previous attempts to model housing markets with agents. Otter et al (2001) focus mainly on locational issues, as do Filatova et al. (2007),

and do not model intermediaries. In England, realtors are important, particularly for their role in setting offer prices. Also, in contrast to Filatova et al., we do not model buyers' and sellers' utilities, except in the simplest possible sense, nor assume utility maximisation. Instead, we just posit that house buyers wish to purchase the most expensive (and presumably, therefore, highest quality) house that they can afford.

Our main objective is to replicate key features of a real world housing market using a simple, stylised agent-based model. In a housing market, prices are:

- 'Sticky' downwards because housing stock is fixed in the short term, so sharp *increases* in demand will increase prices quickly, although these higher prices will in turn draw more houses into the market and eventually ameliorate the price rise. In contrast, sharp *decreases* in demand will will have only a small short-term effect on prices as homeowners will stay put rather than sell, although in the longer term prices will fall.

- Sensitive to the number of newcomers to the market ('first time buyers'). Without newcomers, demand will fall (through migration and death) and so will house prices.

- Sensitive to prevailing interest rates because buyers need to raise mortgages to buy houses: a rise in the mortgage interest rate will reduce the price that owners can afford, *ceteris paribus*.

The model does not simulate any particular geographical location. There is intentionally no attempt to represent either the characteristics of individual housing units, nor spatial attributes such as proximity to facilities. Houses are differentiated only by a neighbourhood index that is based on the prices at which other houses in the locality have been sold.

We first describe the model and then use it to examine the effects of varying policy-relevant parameters from a baseline corresponding to the situation in England in the fall of 2008.

The model

In broad outline, the model simulates a world in which households move either when forced to exit the area (e.g. due to death or a job relocation) or when their mortgages become too expensive for their income, when they try to trade down, or when their incomes rise to the point where they can
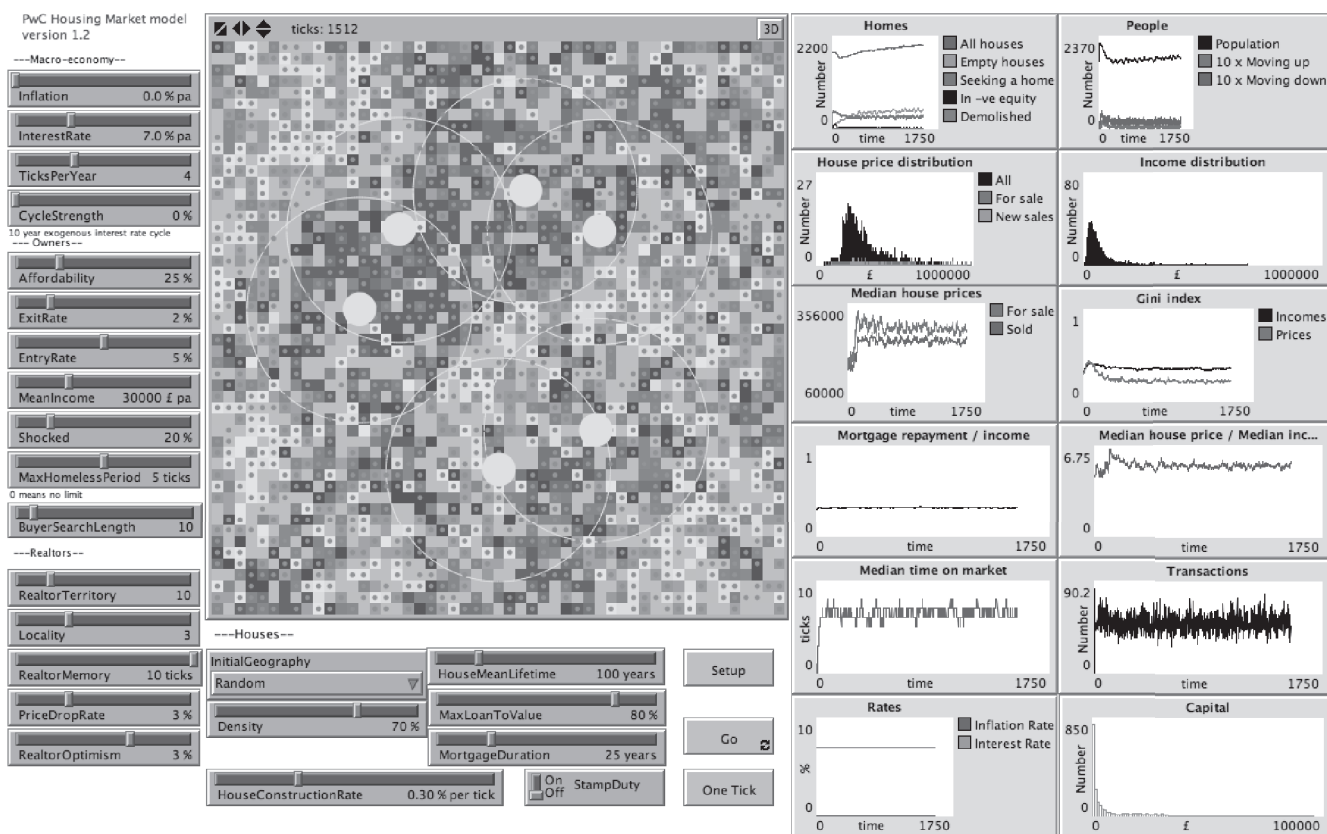
Figure 1: *The user interface of the model, showing the parameters that can be changed, the neighbourhoods that emerge, and the output plots. The original is in colour (contact the first author for a colour version of the figures). The six large light grey (yellow) spots are the realtors. The cells shaded from dark grey (blue) to very light grey (red) are houses coded by their last selling price. Those with grey dots in the centre of the cells are occupied; the others are vacant. Medium grey (green) cells are parkland.*

comfortably afford a more expensive house. Having decided to sell, they approach one or more nearby realtors for valuations, which are based on the realtor's experience of recent price levels in that locality. The seller puts the house on the market at the highest valuation initially but reduces the price later if it does not sell. Buyers look for a house that best fits their income level and make an offer. But offers only go through if chains do not break down, just as in real life. The model works out which chains of offers remain intact and these people then move. Other offers fail as chains collapse. New entrants become discouraged and leave the area after a certain period if they cannot find a property that matches their income.

In more detail (names in `plain font` correspond to parameters of the model that can be varied, see Figure 1):

Grid

There is a 'town' grid of 50×50 plots, of which `Density` have houses on them and $(1 - \text{Density})$ are non-housing (e.g. parks and fields).

Of the houses, some proportion (`InitialVacancy`

`Rate`) are initially vacant and available to buy (at prices set by the highest valuation of the realtors covering the relevant area as explained below). The values of the houses (as indicated by price) is initially distributed either randomly (i.e. with a uniform random distribution), stratified according to their distance from the south-west corner, or randomly clustered, as indicated by `InitialGeography`. Houses also have a quality index, set at the time they are constructed, which is the ratio of the average of the prices of other houses in the locality to this house's sale price.

Sellers

Each period a random `ExitRate` per cent of homeowners exit the 'town' and have to sell at the price set by the realtor as described below. If they cannot sell in the current period, their house remains empty and on the market.

Each homeowner initially has an income selected at random from a Gamma distribution (Salem and Mount 1974) with parameters 1.3 and 5×10^{-5} multiplied by the setting of `MeanIncome`, and a mortgage of value equal to the ratio of `Affordability` to the `InterestRate` multiplied

by the homeowner's income. It is assumed that initial house purchases are mortgage-financed using a repayment mortgage of duration `MortgageDuration` (e.g. 25 years). Borrowers have to find a deposit from their own resources in addition to the mortgage. The amount of deposit is determined by the `MaxLoanToValue`.

Each period `Shocked` per cent of homeowners face an income shock of +20 per cent and the same percentage of homeowners a shock of -20 per cent. This means that their income is suddenly and permanently raised or lowered by this percentage.

Homeowners trade down when the ratio of their mortgage repayment to their income becomes greater than twice the `Affordability` and trade up when the ratio becomes less than half the `Affordability`—they are then 'sellers'.

When they decide to trade, sellers contact all realtors in the area covering their home and ask for valuations. They select the highest valuation, which is the sale price.

Realtors

There are three designated areas of the town, each with a centre and with overlapping radiuses of coverage together encompassing most of the town. Realtors are located in the centre of each area, two to an area, to represent clustering. Most houses will therefore be covered by at least two realtors. Realtors have a radius of coverage given by `RealtorTerritory`. Houses not in a realtor's territory take their valuations from the nearest realtor.

Realtors know the prices of all houses they have recently sold. At the beginning of the simulation run, the mortgage value of each house is lodged with one local realtor, to give realtors an initial starting point for their valuations.

When asked for a valuation after the first period, the realtors use the median prices of houses that they have sold in the last `RealtorMemory` periods that are within a certain radius of the house for sale (the `Locality`) multiplied by the house's quality index. If a realtor has no sales in the locality, it uses the median price of all sales in the `Locality` (whenever they were made) or if there are no sales, the average price of all houses in its territory. The valuation is boosted by a `RealtorOptimism` factor (e.g. 3 per cent), which allows for realtors hoping that they can sell for more than the going rate.

If a house fails to sell during a period, the selling price is reduced by `PriceDropRate` per cent and it continues to be on the market for the next period and until it is sold or demolished.

Buyers

Each period there are new entrants to the town equalling `EntryRate` per cent of the existing population of homeowners.

The new entrants and non-exiting sellers search randomly across the whole town for `BuyerSearchLength` properties that are for sale, for which no offer has yet been made, and that they can afford (i.e., those for which their mortgage repayment would not be greater than the `Affordability` for their income and for which they would have a sufficient

deposit to hand). They can also put any capital that they have gained from selling a previous house towards the cost of their desired property. They then choose the one property that comes closest to their maximum affordable price, and make an offer at the price set by the seller. Thus the first buyer to make an offer on a property has their offer accepted.

If an entrant has failed to purchase a house for `MaxHomelessPeriod` periods after entry, it exits.

Transactions

An offer only turns into a successful sale if the house being bought is either already empty, or the seller succeeds in moving to another house that they have purchased. Hence, transactions can fail if the chain cannot complete.

The exiters move out of their houses and potential buyers (if any) that made successful offers for the now empty houses move in. The purchasers' houses are then moved into by those who had made offers for them and so on down the chain. Once all possible transactions are complete, the next time step starts.

When a house is sold, the seller gets the sale price and immediately uses as much of this as necessary to pay off any remaining mortgage. If money remains after paying off the mortgage, this capital can be used as partial or full payment for the house the seller is moving into. If the sale price is insufficient to pay off the mortgage, the seller is in a position of 'negative equity' and has to withdraw the house from the market.

The realtor concerned records details of successful sales and uses these records when valuing houses within the same `Locality`.

If an offer fails to go through in one period, it immediately lapses.

New build and demolition

At each time step, new houses are added to the town at random empty locations, at a rate of `HouseConstructionRate` per cent of the existing housing stock, unless there are no empty plots. A house has a lifetime set at the time of its construction, drawn from a random exponential distribution with a mean of `HouseMeanLifeTime`. At the end of its life, it is demolished and the plot becomes vacant again and available for house building. A house is also demolished if its sale price falls below one tenth of the median price of all houses. If there is someone living in a house being demolished, they attempt to buy another house and if they have not succeeded after `MaxHomelessPeriod` timesteps they leave the town.

Macroeconomic variables

As well as the (mortgage) `InterestRate`, a rate for (wage) `Inflation` can be set. Both these are measured in per cent per year. The time represented in the model can be calibrated in `TicksPerYear` (i.e. timesteps per year). The interest rate may be periodically varied (along a sine wave) by varying `CycleStrength`.

Display

Houses are assigned a colour taken from one of ten shades from blue to red according to their price (based on last transaction, or mortgage value initially), see Figure 1.

There are time plots of:

- The total number of houses, the number that are empty, and the number that are being demolished at each time step

- The number of people without a house who are seeking one and the number in a house, but with negative equity

- The total number of people in the town, and the numbers who are moving because the ratio of their mortgage repayments to their income is greater than twice or less than half the `Affordability` ratio

- Median house prices and median incomes (and the ratio between the two)

- Measures of income and house price inequality (Gini co-efficients)

- Transaction volumes

- The median time that a house for sale is on the market before it is sold or withdrawn from sale.

For the current time period, there are histograms for

- the distribution of incomes and house prices

- the distribution of capital accumulated by owners selling houses.

The behaviour of the model

The model is written in NetLogo (Wilensky 1999) and may be run as an applet on the web at `http://cress.soc.surrey.ac.uk/housingmarket/ukhm.html`. Our objective was to reproduce the main characteristics of the market as listed above. This has been achieved. Some examples[1] are described below. The examples are based on model parameters that have been calibrated to match approximately the actual English situation in the early fall (autumn) of 2008.

1. The price elasticity with respect to the prevailing interest rate is -1.03 (Figure 2). Thus a 43 percent increase in the interest rate (from 7 to 10 percent) results the median house price falling by 47 per cent, with an immediate sharp decrease in the number of transactions and the ratio of mean house prices to mean income reduced from about 5.5 to about 3.6. This is the result of buyers finding that they cannot afford more expensive houses, driving down demand and thus the market price.

2. First time buyers compensate for the drain of owners who leave due to migration or death. If prices are too high, the newcomers cannot find a property to buy and eventually leave, discouraged, which drives down prices (for example, if the newcomers have an income 20 per cent less than the incumbents, £24,000 compared with £30,000, the median house price drops by about the same ratio, see Figure 3).

[1]Although the model is stochastic, the variations between runs are small compared with the qualitative effects described below.
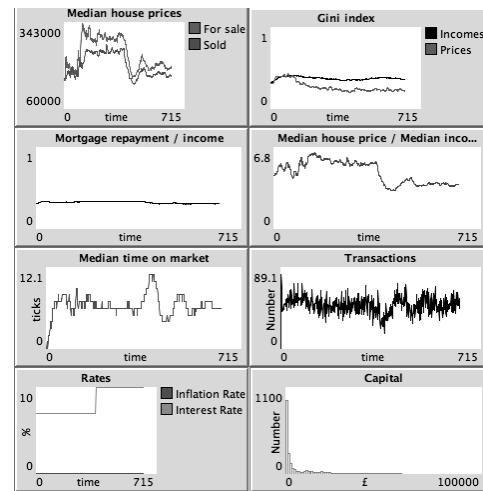

Figure 2: *The effect of an abrupt change in the mortgage interest rate at time step 360 on a range of market features*
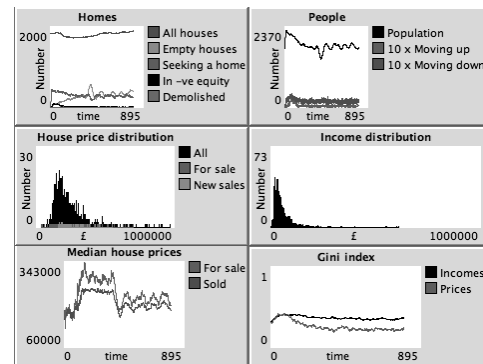

Figure 3: *The effect of an abrupt 20 per cent decrease in the income of newcomers at time step 360 on a range of market features*

3. A large demand shock (e.g. a surge of 500 newcomers into a market of 2,200 houses) produces an immediate general increase in house prices of about 30 per cent and a near doubling of the price to income ratio, which slowly returns to normal as more houses come onto the market.

Conclusions from the model

By starting the simulation with a 'baseline' set of parameters and then varying one of the parameters while the simulation is running, we can examine the sensitivity of indicators such as the median house price to exogenous economic and policy changes. A baseline parameter set of:

1. mortgage interest rate: 7%

2. maximum loan to value ratio: 100%

3. new entrants to the town have the same distribution of income as the residents

4. housebuilding is at the rate of 0.33% of the housing stock in each period
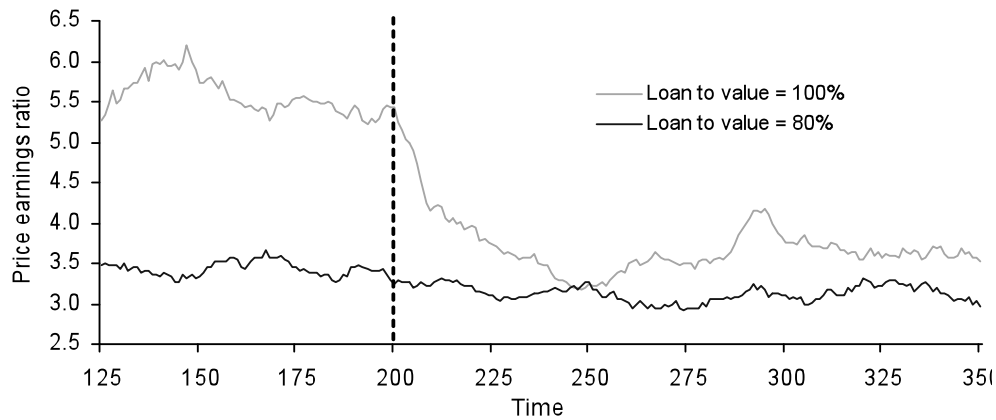
Figure 4: *House price to earnings ratio in response to an interest rate increase at period 200 from 7% to 10%*

was used. With this baseline model, we see that both house prices and the number of housing transactions tend to be volatile in the short term (see Figure 2, helping to explain why short-term housing market forecasts are often subject to significant margins of error.

In the longer term, average house prices tend to show a stable relationship to average income levels for given levels of mortgage interest rates and loan-to-value ratios; if interest rates fall or maximum loan-to-value ratios rise, then the equilibrium house price to income ratio rises in response. This can be seen in Figure 4. The upper line shows the effect of a change in interest rate from 7% to 10% with a loan-to-value ratio of 100%. As expected, the rate change induces a sharp drop in prices. However, if we now alter the maximum loan to value ratio to 80%, the effect of an interest rate change from 7% to 10% is modest (see Figure 4). That is, the interest rate effect on house prices is much weaker when the maximum loan-to-value ratio is lower; in particular, the model suggests that moving from a 100% to an 80% maximum loan-to-value ratio can have a major depressing impact on house prices through keeping first-time buyers with limited savings out of the housing market and this can outweigh any positive effects on house prices and transactions volumes from interest rate cuts.

With an 80% maximum loan-to-value, new entrants to the market are mostly constrained by their ability to find a deposit of at least 20% and only secondarily by their ability to afford interest repayments. Interest rate rises do still have some downward in uence on house prices, but it is much less marked in this case as the lower line in Figure 4 shows. Hence, if, as has happened during 2008, maximum loan-to-value ratios are reduced from 100% to, say, around 80% then this will not only reduce house prices significantly but may also reduce the significance of mortgage interest rate cuts in potentially reviving the market—if first-time buyers (and other households with limited savings) cannot afford a deposit, then the mortgage interest rate does not matter much to them.

Other interesting conclusions that can be drawn from the model include:

- Even when starting a run with a 'random' `InitialGeography`, where house quality and owners are randomly distributed across the town, clusters or 'neighbourhoods' of high priced and low priced houses tend to emerge (see Figure 1 for an example). This is because high priced houses in a locality raise the prices of other houses around them by in uencing the valuations given by realtors in the neighbourhood. These relatively expensive houses attract relatively rich owners, creating a well-off neighbourhood. The reverse is true for poor neighbourhoods. The positions of the clusters is entirely random and varies for each run. By setting the `InitialGeography` to 'clustered', a starting configuration with neighbourhoods can be created for the simulation and thee neighbourhoods are maintained, more or less, as the simulation runs. If certain areas had natural advantages (such as pleasant surroundings or easy access to facilities), one would expect neighbourhoods would crystallise around them and this is likely to be the way real neighbourhoods form and persist, changing only slightly across centuries.

- Adding wage in ation does not change the qualitative behaviour of the model: all prices rise in parallel with the in ation, but the shape of the house price distribution and the Gini indices do not change.

- Adding the British stamp duty land tax (a tax levied on house sales, which rises as a percentage as the price of the property increases) to the simulation shows that altering the rate of duty has a barely detectable effect on house prices. If, for example, stamp duty is set at 1% rather than zero then the impact is effectively lost in the short-term noise seen in house prices. The implication is that the temporary holiday on the 1% stamp duty rate for houses between £125,000 and £175,000 that was announced in early September 2008 seems unlikely to have had more than a marginal impact on English house prices

compared with other market drivers like actual and expected household income levels, mortgage interest rates and credit availability.

Much the same applies to the impact of varying other fixed costs of moving, such as realtor's fees and removal costs.

- Shocks to the market can generate long-lasting oscillations, giving the effect of emergent cycles of boom and bust. An example can be seen in Figure 2 after the change in interest rate at step 360.

Developments

The aim of the model is to investigate the complicated interactions that create a housing market, and so it is desirable for the model to be as simple as possible, while showing the required emergent behaviour. There are a number of features of real housing markets yet to be studied: some of them may prove to be important and others less so. These include:

- the rental sector

- repossessions and forced sales

- house building that responds to demand, but with a lag for construction

- endogenising the mortgage interest rate, so that it responds to demand

- the effects of competition between realtors: those more successful in selling houses would benefit from increased profits and be able to increase their radius of coverage.

It may also be necessary to reproduce a real location in order to validate the model and to account for spatial effects such as distance to urban centres and jobs. Nevertheless, we have demonstrated that the agent-based modelling of a housing market is both possible and potentially useful.

Acknowledgement

References

Filatova, T.; van der Veen, A.; and Parker, D. C. 2007. Modeling of a residential land market with a spatially explicit agent-based land market model (ALMA). In *2nd Workshop of Market Dynamics SIG*.

Otter, H. S.; van der Veen, A.; and de Vriend, H. J. 2001. ABLOoM: Location behaviour, spatial patterns, and agent-based modelling. *J. of Artificial Societies and Social Simulation* 4(4):2. http://jasss.soc.surrey.ac.uk/4/4/2.html.

PricewaterhouseCoopers. 2007. UK economic outlook November 2007. Technical report. http://www.pwc.co.uk/pdf/uk_economic_outlook_sec4_nov07.pdf.

Salem, A. B. Z., and Mount, T. D. 1974. A convenient descriptive model of income description: the gamma density. *Econometrica* 42:1115–1127.

Wilensky, U. 1999. Netlogo. http://ccl.northwestern.edu/netlogo/.