Computers, Environment and Urban Systems 84 (2020) 101541

Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

journal homepage: www.elsevier.com/locate/ceus

Decentralized markets and the emergence of housing wealth inequality

Omar A. Guerrero

Department of Economics, University College London, The Alan Turing Institute, UK

ARTICLE INFO

Keywords:
Housing
Inequality
Agent-based model
Life cycle
Micro-data

ABSTRACT

Recent studies suggest that the traditional determinants of housing wealth are insufficient to explain its current inequality levels. Thus, they argue that efforts should focus on understanding institutional factors. From the perspective of complex adaptive systems, institutions are more than the 'the rules of the game', they also consider the interaction protocols or the 'algorithm' through which agents engage in socioeconomic activities. By viewing markets as complex adaptive systems, I develop a model that allows estimating how much housing wealth inequality is attributable to the market institution. It combines virtues from two different modeling traditions: (1) the microeconomic foundations from overlapping-generation models and (2) the explicit interaction protocols of agent-based models. Overall, the model generates prices and housing inequality endogenously and from bottom-up; without needing to impose assumptions about the aggregate behavior of the market (such as market equilibrium). It accounts for economic and institutional factors that are important to housing consumption decisions (e.g., wages, consumption of goods, non-labor income, government transfers, taxes, etc.). I calibrate the model with the British Wealth and Assets Survey at the level of each individual household (i.e., ~25 million agents). By performing counter-factual simulations that control for data heterogeneity, I estimate that, in the United Kingdom, the decentralized protocol interaction of the housing market contributes with one to two thirds of the Gini coefficient. I perform policy experiments and compare the outcomes between an expansion in the housing stock, a sales tax, and an inheritance tax. The results raise concerns about the limitations of traditional policies and call for a careful re-examination of housing wealth inequality.

1. Introduction

In recent years, housing has become central in the broader discussion of wealth inequality (Allegre & Timbeau, 2015; Bonnet, Bono, Chapelle, Wasmer, et al., 2014; Piketty, 2014; Rognlie, 2016; Stiglitz, 2015).[1] Among housing scholars, however, the problem of housing wealth inequality (HWI) has been an important topic for quite some time (Appleyard & Rowlingson, 2010; Arundel, 2017; Dewilde, 2011; Doling & Ronald, 2010; Forrest, 1995; Forrest, Murie, & Williams, 1990; Henley, 1998; Lersch & Dewilde, 2018; Maxwell & Sodha, 2006; Ronald, Kadi, & Lennartz, 2015; Rowlingson, 2002; Ryan–Collins, 2018; Ryan-Collins, Lloyd, & Macfarlane, 2017; Wind, Lersch, & Dewilde, 2016). In the United Kingdom, public debate around HWI has become part of the usual content of media outlets (Bulman, 2018; Pettifor, 2018; Weaver, 2017). Think tanks, journalists, politicians, academics and, more broadly, the civil society constantly engage in these discussions, proposing different remedies (Bowie, 2017; Harris, 2018; Library, 2017; Pettifor, 2018). These diagnostics are often persuasive and enjoy of a coherent theoretical backbone. Yet much of the supporting evidence seems rather descriptive. Among those providing causal evidence on the determinants of homeownership (Andrews, 2010; Andrews & Sánchez, 2011; Atterhög, 2005; Fisher & Jaffe, 2003; OECD, 2011) (mainly through regression analysis), some have pointed out that traditional structural factors such as class, age, income, and education are insufficient to explain today's HWI; suggesting that efforts should be directed towards understanding the role of institutions (Arundel, 2017; Dewilde, 2011; Wind et al., 2016).

In a seminal study, Dewilde (2011) uncovers severe methodological and data-related limitations in studying the determinants of HWI. Furthermore, she makes the case for the importance of institutional context specificity. That is, cross-national differences of institutional drivers may interact with various social mechanisms in (non-trivial) ways that are specific to each country. Thus, within-country institutional factors should be carefully analyzed in order to understand HWI. While Dewilde concentrates on the role of the welfare estate, an intriguing idea can be derived from her argument: *institutions cause HWI*. Allow me to take this idea one step further and focus on a specific institution that prevails in most economies, and with various degrees of

E-mail address: oguerrero@turing.ac.uk.

[1] A recent paper by Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019) studies the performance of a large class of assets and finds housing to have remarkable long-term returns.

https://doi.org/10.1016/j.compenvurbsys.2020.101541
Received 17 February 2020; Received in revised form 6 August 2020; Accepted 8 August 2020
Available online 26 August 2020
0198-9715/

intervention: *the market*. That the decentralized protocol, algorithm, or *tatonement* process behind buying and selling properties generates skewed distributions of housing wealth calls for a profound rethinking on the data and methods typically used to study housing dynamics. Furthermore, it carries major implications regarding the limitations of traditional redistributive instruments because, through the structure of its interactions, the market may restrict the potential outcomes of policies.

In order to conceptualize the housing market, I adopt a definition by Axtell (2005), which offers a procedural view that is more in line with the real world[2]:

> "*... a heterogeneous population of autonomous entities, each of whom has internal states that describe its self-interest as well as certain external states. Each entity is engaged in purposive activity to further its interests... Each individual receives information from other individuals directly, and has access to some global state information as well, although no agent has complete information on the global state.*" (ellipsis supplied).

This view conceives the market as a complex adaptive system, which Gatti, Gallegati, and Kirman (2000) summarize in three fundamental issues: (1) the heterogeneity of the agents in the economy, (2) the ways in which agents interact and (3) the dynamic process which governs the evolution of the individual and the aggregate variables. In this paper, heterogeneity comes from empirical household data, interactions are assumed to be decentralized pairwise negotiations, and dynamics are produced through through the agent's interdependent trajectories of homeownership.

Empirically evaluating market effects on HWI represents a challenge for which traditional statistical tools are not designed. Most available datasets do not allow identifying the market as something that can be controlled for (consistent with Dewilde's argument about limited data). For instance, a cross-national panel approach is ill-suited as it is impossible to find comprehensive data where a clear institutional counterfactual to the market exists. Likewise, natural experiments where a market has been entirely replaced by a different institution are nonexistent, at least in any useful data form. Moreover, any study evaluating the impact of a specific intervention would only be able to address exactly that, the intervention.

In this paper, I attempt to provide first estimates on how much HWI is caused by the market. My approach proposes an agent-based model. A computational approach is ideal for this problem because it allows creating synthetic populations. Of course, for this to be empirically relevant, the model should be parsimonious enough to be calibrated with real-world data, trying to minimize any overfitting issues. In addition, it should be granular in order to account for all available sources of heterogeneity existing in empirical micro-data. The model presented in this study has those attributes. The rest of the paper is organized as follows. In the reminder of this section, I explain the general philosophy of the model and summarize some of the main findings. Section 2 presents the microeconomic foundations of the model. Section 3 shows how to implement it computationally. In Section 4, I demonstrate its empirical application. Section 5 provides results on three policy experiments. Finally, I offer some thoughts on future directions and conclusions in Section 6.

[2] Traditional economic models typically assume no interactions, as agents respond to a price vector that does not emerge from bottom-up, but that is produced through aggregate equilibrium assumptions adopted for mathematical convenience. While there are game-theoretic models that try to deal with agent-to-agent strategic interactions, unrealistic assumptions about rationality (also motivated by mathematical convenience) render them inflexible to consider realistic institutional settings. Thus, the approach 'agent-based artificial markets' is preferred to study the problem at hand.

1.1. Modeling approach and main findings

My approach consists of simulating all 25 million households in Great Britain, allowing them to buy and sell real estate in a decentralized fashion through agent-computing or agent-based modeling. An overlapping-generations microeconomic foundation sits at the core of this model (Kraft, Munk, & Wagner, 2018; Yang, 2009). It considers finitely-lived agents who make choices about labor, consumption, and housing. A distinctive feature that departs from traditional overlapping-generations models is that, here, housing decisions are 'asynchronous' with respect to labor/consumption ones. This stems from the fact that housing transactions happen at a different time scale.[3] Such asynchronicity allows the model to generate HWI endogenously, even in homogeneous populations.

The main finding is that the market is responsible for one to two thirds of the Gini coefficient of HWI in Great Britain (it varies depending on the region). I also perform three policy experiments: an expansion of the housing stock, a sales tax, and an inheritance tax. While the outcomes of these experiments should not bare any prescriptive weight (since they are considerably stylized), they point out interesting differences on how sensitive the different regions could be to such interventions.

1.2. Related literature

Agent-based modeling in housing studies is not new. Pioneering papers date back to the early 2000's (Torrens, 2001), while landmark models were developed later to understand land-use problems (Filatova, Parker, & Veen, 2009; Parker & Filatova, 2008). During the global financial crisis of 2008, agent-computing became a popular choice to model housing dynamics (Gilbert, Hawksworth, & Swinney, 2009; McMahon, Berea, & Osman, 2009) and housing bubbles (Baptista et al., 2016; Dieci & Westerhoff, 2012; Erlingsson et al., 2014; Erlingsson, Raberto, Stefánsson, & Sturluson, 2013; Ge, 2014; Ge, 2017; Geanakoplos et al., 2012; Kouwenberg & Zwinkels, 2015). In addition, agent-based models have also been deployed to understand residential segregation (Feitosa, Le, & Vlek, 2011; Jordan, Birkin, & Evans, 2012; Pangallo, Nadal, & Vignes, 2019; Yin, 2009).

When looking at this literature, price formation seems to be an area of disagreement. Some models consider exogenous prices. Others generate endogenous prices through aggregate supply and demand. Some other models emerge prices by specifying auctions between the seller and multiple potential buyers.

On the side of more traditional economics, equilibrium models dealing with household dynamics tend to be highly stylized and rigid (in order to obtain mathematical solutions). However, their parsimony provides intuitive micro-foundations and facilitates calibration with real-world data. This paper combines features of both agent-based and overlapping-generations modeling traditions.

Before proceeding, it is important to acknowledge that this is a first attempt to develop such a model and, as with every model, it is an abstraction of reality. The purpose of this particular abstraction is not only to provide theoretical insights, but also to generate a tool that can be used with existing micro-data; thus, it requires to be parsimonious. Since the goal is to provide a mechanism through which prices and HWI emerge endogenously, the paper does not address spatial structure, financial instruments, and rental markets. These factors, however, may have important quantitative implications. For instance, spatial structure can restrict interactions. Financial instruments may expand or contract the agents' budgetary constraints. Rental markets could offer a side

[3] Rowlingson (2002) provides evidence on lack of long-term financial planning by households and makes a compelling argument for its relevance to housing choices and for the inadequacy of overlapping generations-style models.

option to buying housing, as well as potentially incentivize rent-seeking behavior. And all of them could increase or decrease HWI in a model. While these quantitative implications are, indeed, important to provide more reliable estimates, they also demand further modeling assumptions and additional free parameters that need to be calibrated. Therefore, given the data available for this study, I will consider these factors in future extensions.[4]

2. Microeconomic foundations

I assume a fixed population of N finitely-lived agents and a constant amount of real estate in the economy; i.e. a fixed stock. One can think of housing as an asset that can be infinitely divided, so that any agent may hold a share. This context can be interpreted as a short-term scenario or as an economy with a stagnant housing supply. Given that the common asset cannot be created or destroyed, the only way to increase its ownership is by buying shares from other agents. In this model, purchases of the common asset take place in a decentralized fashion, resulting from pairwise interactions.

Agents maximize utility over consumption, labor, and housing. Housing decisions, however, take place at a different moment because of discrepancies in timescales with regard to consumption/labor decisions. For instance, when a person engages in a housing transaction, the horizon for the next transaction (if any) is considerably longer than for the next decision about, for example, paying the electricity bill. However, should another real estate opportunity appear, the agent revises their consumption/labor lifetime plan, taking into account the possibility of updating their share of the common asset. This loose relationship between consumption/labor choices and housing decisions generates sub-optimal transactions. It turns out that sub-optimal transactions are one of the sources of endogenous inequality in the model. Next, I formalize these ideas.

2.1. The agent

Every period, agent i consumes c_i and enjoys of ℓ_i time-units of leisure. This utility is 'enhanced' through ownership of the common asset. The degree of enhancement depends on (1) the amount A_i of asset owned and (2) the agent's preferences β_i towards hosing.[5] Every period, agent i survives with probability δ. When an agent dies, they are replaced by an identical (but younger) agent who inherits A_i. The lifetime utility stream of agent i is given by

$$U_i = \sum_{t=0}^{\infty} \gamma^t (1 + \beta_i A_i) U_{i,t}(c_{i,t}, \ell_{i,t}), \tag{1}$$

where γ is the discount rate.

The consumption/labor maximization problem follows the textbook microeconomic model. For tractability, I do not consider savings. Thus, in absence of inter-temporal choices, the agent's problem simplifies into

$$\max_{c_i, \ell_i} \quad U_i = \left(\frac{\gamma^{L_i}}{1-\gamma}\right)(1 + \beta_i A_i) c_i^{\alpha_i} \ell_i^{1-\alpha_i}$$
$$\text{s. t.} \quad c_i = (1 - \ell_i)\tau_i w_i + z_i B_i + s_i, \tag{2}$$

where w_i is the real wage per time unit, B_i is a constant amount of non-labor income (such as financial dividends) received every period, $1 - \tau_i$ is the labor-income tax rate, $1 - z_i$ is the non-labor-income tax rate and s_i are government transfers. Exponent L_i represents the age of the agent.

The optimal level of utility is

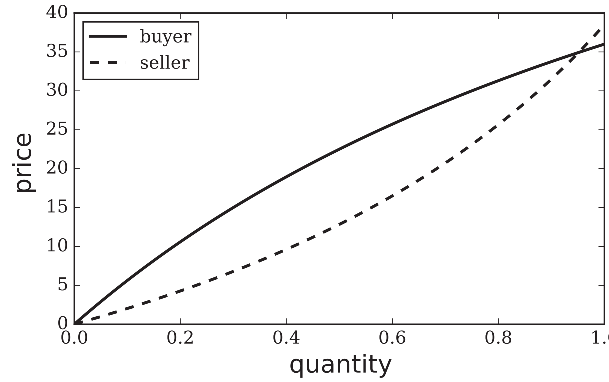

Fig. 1. Transaction space.

$$U_i^* = \mu_{L_i}(1 + \beta_i A_i)(\tau_i w_i + z_i B_i + s_i)\widehat{U_i}, \tag{3}$$

where $\mu_{L_i} = \frac{\gamma^{L_i}}{1-\gamma}$ and $\widehat{U_i} = \alpha_i^{\alpha_i}\left(\frac{1-\alpha_i}{\tau_i w_i}\right)^{1-\alpha_i}$ for compactness.

2.2. Transactions

Transactions take place in a pairwise fashion. The encounters happen at random, with equal probability, and independently of each other.[6] If a sale is successful, the agents agree on transacting a quantity q of the common asset for a price p. Whether a transaction is successful or not depends on incentive compatibility.

Consider a purchase where buyer i and seller j face p and q. The buyer's utility surplus from this transaction is

$$\begin{aligned} S_i = {} & [1 + \beta_i(A_i + q)](\tau_i w_i + z_i B_i + s_i - p)\widehat{U_i} + \\ & (\mu_{L_i} - 1)[1 + \beta_i(A_i + q)](\tau_i w_i + z_i B_i + s_i)\widehat{U_i} - \\ & \mu_{L_i}(1 + \beta_i A_i)(\tau_i w_i + z_i B_i + s_i)\widehat{U_i}. \end{aligned} \tag{4}$$

Eq. (4) captures the difference between the buyer's utility from going through with the transaction and the utility from not doing it. The top line represents the utility of the buyer in the current period, where they would pay p and start enjoying of the additional asset q right away. The middle line is the utility from the transaction for the rest of the periods. Because purchases require a one-off payment, the agent enjoys q for the rest of their life without making any further payments to the seller. Finally, the bottom line is the present-valued utility stream that the buyer would receive without the transaction. Therefore, if $S_i > 0$, then the buyer has incentives to acquire the proposed q at the proposed price p.

The surplus equation for the seller can be constructed in a similar way as

$$\begin{aligned} S_j = {} & [1 + \beta_i(A_i - q)](\tau_j w_j + z_j B_j + s_j + hp)\widehat{U_j} + \\ & (\mu_{L_j} - 1)[1 + \beta_i(A_i - q)](\tau_j w_j + z_j B_j + s_j)\widehat{U_j} - \\ & \mu_{L_j}(1 + \beta_i A_i)(\tau_j w_j + z_j B_j + s_j)\widehat{U_j}, \end{aligned} \tag{5}$$

where $1 - h$ is the tax rate for selling the common asset.

In order for a proposed transaction to be feasible, it must be the case that $S_i > 0$ and $S_j > 0$. Under these conditions, there can be a multiplicity of possible agreement between the two agents. These potential transactions live in the space between the agents' indifference curves. By setting $S_i = 0$ and $S_j = 0$, it is possible to construct these curves as

$$p_A^i = \frac{\mu_{L_i}(\tau_i w_i + z_i B_i + s_i)\beta_i q}{1 + \beta_i(A_i + q)}, \tag{6}$$

and

[4] Note, however, that most existing models also lack these one or more of these factors.

[5] Note that β_i is not easy to obtain from data. Should the author not be comfortable with the estimation strategy presented in this paper, I suggest setting $\beta_i = 1$ since the model becomes simpler and the qualitative features of the results remain broadly unaltered.

[6] Spatial considerations would modify this assumption.

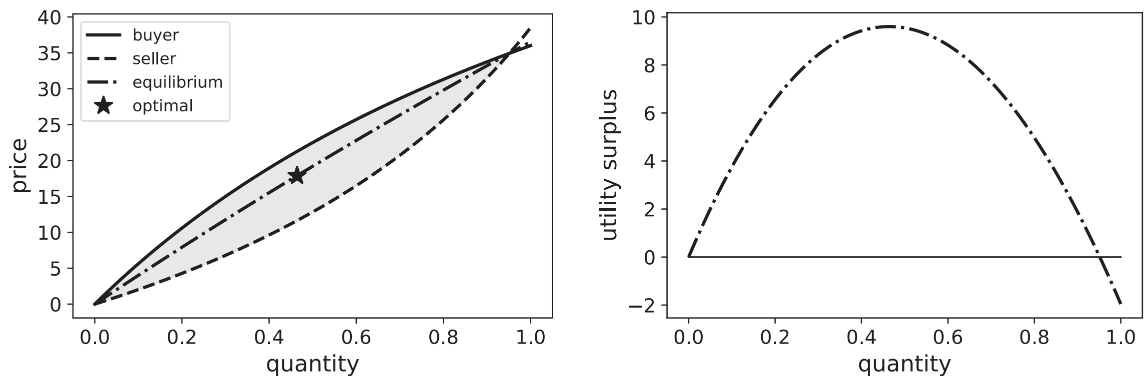

Fig. 2. Equilibrium and optimum.

$$p_A^j = \frac{\mu_{L_j}(\tau_j w_j + z_j B_j + s_j)\beta_j q}{h_j[1 + \beta_i(A_i - q)]}. \tag{7}$$

Fig. 1 shows the buyer's and seller's indifference curves in the q–p space. On the one hand, the buyer is willing to trade at any price at or bellow the solid line. On the other, the seller will only accept offers at or above the dashed line. Any transaction between these curves is feasible.

There may be several mechanisms that could lead to a particular point in the transaction space (or even outside, e.g. in the case of erroneous utility estimation). For computational efficiency, I focus on one particular outcome that can be obtained without an explicit bargaining process.

2.3. Pairwise equilibrium

Suppose that the transaction space between two agents in a given period exists. Then, for some quantity q, any corresponding price in the transaction space would yield a feasible transactions. Let p^e denote the equilibrium price for q, such that $S_i = S_j$. That is, the equilibrium price for a given quantity is the one in which both agents obtain the same utility surplus (although their total utilities may differ). It follows that, for any quantity inside the transaction space, there is a unique equilibrium price. This determines the equilibrium path in the transaction space (see left panel in Fig. 2), which is described by

$$p^e = \frac{[\mu_{L_i}\beta_i\widehat{U}_i(\tau_i w_i + z_i B_i + s_i) + \mu_{L_j}\beta_j\widehat{U}_j(\tau_j w_j + z_j B_j + s_j)]q}{\widehat{U}_i[1 + \beta_i(A_i + q)] + h\widehat{U}_j[1 + \beta_j(A_j - q)]}. \tag{8}$$

2.4. Optimal transaction

For a given quantity, Eq. 2.3 provides a unique equilibrium price. Then, what should be the equilibrium quantity to be chosen? Agents choose the pair (q^*,p^*) of equilibrium quantity and price that maximizes their utility surpluses.

If the transaction space exists, an optimal pair (p^*,q^*) always exists. To show this, one can substitute the price variable p in Eq. (4) (or in 5) with the equilibrium price obtained in Eq. 2.3. The resulting equation is the utility surplus from pairwise equilibrium pricing; and it is a function of q. It turns out that such function is concave in q, as shown in the right panel of Fig. 2 (the explicit function can be found in Appendix A). Thus, by differentiating this function and solving for q, one can obtain the optimal q^*. Then, by imputing q^* back into Eq. 2.3, one obtains the optimal price p^* (explicit solutions are provided in Appendix A).

At this point, it is important to clarify that pairwise equilibrium is a micro-level outcome. Therefore, there are no assumptions about market equilibrium. As I will show, the model produces steady-state outcomes in the behavior of HWI. Nevertheless, this does not imply macro-level equilibrium. In fact, my choice to determine prices through pairwise equilibrium is motivated by computational efficiency (since one does not require a bargaining algorithm). Future versions of this model should implement more realistic price formation mechanisms, perhaps according to the attributes of the specific housing unit on sale; which obviously cannot be currently implemented due to the continuous nature of the common asset.[7]

2.5. Transaction outcomes

Because of the asynchrony between housing decisions and consumption/labor choices, the existence of an optimal transaction does not guarantee its feasibility. This is so because the evolution of A_i is not considered as an inter-temporal utility maximization problem in Eq. (3). Consequently, the budget constraint is not binding in housing decisions. If the budget was fully binding, then optimality would always be feasible. Full bindness is typically achieved through strong rationality assumptions (the so-called time-consistency). This model does not impose such presumptions. In fact, loose bindness is a feature that enables endogenous HWI through sub-optimal transactions.[8]

In total, there are five transactions cases that can take place, all illustrated in Fig. 3. In case 1, the optimal equilibrium is feasible because the buyer has enough income (horizontal dotted line) and the seller owns enough common asset (vertical dotted line). In case 2, the seller owns less than q^*, so they sell everything to the buyer. Case 3 is similar, but the buyer does not have enough income. In case 4, the buyer spends all their income in a sub-optimal fraction of the seller's asset. A similar outcome occurs in case 5, where the seller has enough asset to reach the optimal sale, but the buyer cannot afford it.

Given the discrete nature of the five transaction outcomes, I implement the model computationally and demonstrate its capability to emerge HWI from agent-level interactions.

3. Computational implementation

Allow me to define

$$A = \sum_i^N A_i \tag{9}$$

as the relative abundance of the common asset in the economy. A is the only free parameter to be calibrated (all others come from data). As I show below, variation in the relative abundance generates different

[7] Appendix E shows a generalization of the model in which multiple matching rounds are allowed. Thus, agents pick the best transaction from multiple encounters within the same simulation step. The analysis shows that the main qualitative features of the model remain roughly unchanged. Therefore, I proceed with the single-encounter approach since it is computationally more efficient.

[8] Appendix B shows that, in homogeneous populations, HWI cannot emerge endogenously when optimal transactions are the only possible outcome.

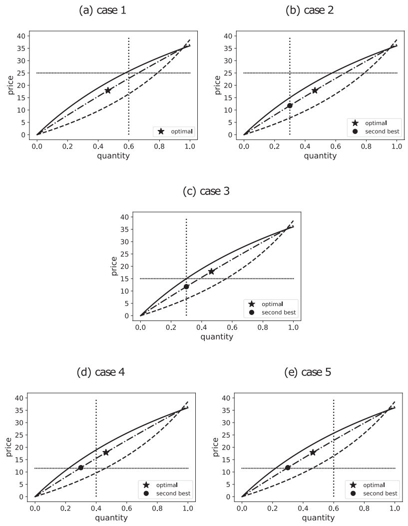

Fig. 3. Transaction outcomes.
The horizontal and vertical dashed lines indicate the buyer's budget constraint and the seller's share of the common asset respectively.

levels of HWI. Algorithm 1 provides the model pseudocode.

Algorithm 1: Agent-computing model.

```
Input: {N, A, B_i, w_i, τ_i, z_i, s_i, α_i, β_i, δ, γ, L_i} for every agent i
foreach time step t do
    foreach agent i do
        increase age L_i;
        survive with probability δ;
        if agent dies then
            reset age;
    match agents;
    foreach match do
        if purchasing space exist then
            compute (q, p);
            update shares A_i and A_j;
```

3.1. Proof of concept

In this section, I present a set of hypothetical simulations. Table 1 enlists all the model parameters and their hypothetical benchmark values. I assume a population size $N = 10{,}000$. Each simulation is run until the system reaches a steady state. Agents are homogeneous and, at $t = 0$, they are endowed with the same amount of the common asset.

To quantify HWI, I employ the Gini coefficient

$$\mathscr{G} = \frac{2\sum_{i}^{N} A_i}{N\sum_{i}^{N} A_i} - \frac{N+1}{N}, \tag{10}$$

where $A_i \leq A_{i+1}$ for every i.[9]

Panel (a) in Fig. 4 shows individual trajectories of A_i throughout a simulation. These endogenous dynamics are decomposed into the five transaction cases presented in panel (b). Then, panel (c) shows that the Gini coefficient achieves steady-state behavior.

[9] I measure inequality in the number of properties owned and not in their values because the model does not have an explicit representation for a housing unit. That is, value is an attribute of a specific property, and this is ill-defined in the current representation of the common asset (an infinitesimally divisible good); something to improve in a future revision.

Table 1
Model parameters arranged by source.

Parameter	Concept	Benchmark Value
A	Relative abundance	10,000
N	Number of agents	10,000
α_i	Consumption/labor preferences	0.500
β_i	Housing preferences	0.500
γ	Discount rate	0.950
δ	Survival probability	0.950
w_i	Wage	1.000
B_i	Non-labor income	0.000
s_i	Government transfers	0.000
$1-\tau_i$	Labor income tax	0.000
$1-z_i$	Non-labor income tax	0.000
$1-h$	Sales tax on common asset	0.000

There are a couple important insights to be obtained from Fig. 4. First, the model is capable of generating HWI endogenously, even when the population is homogeneous. Here, the aging process is the stochastic element that triggers a baseline level of heterogeneity, enabling transactions between otherwise identical agents. Section 3.2 explains in detail the role of the aging process and how it is possible to control for it. Second, the aggregate level of inequality is stable, and could mislead the reader into thinking that there is an aggregate equilibrium. Recall that the model establishes micro-level pairwise equilibria in the transactions, but does not assume anything at the population level. In fact, the diverse trajectories in panel (a) suggest that, instead, the housing marker is in constant disequilibrium. Thus, assuming system-wide equilibria just because one observes certain stability in aggregate data can lead to misspecifications in which institutions are replaced by fictitious coordinating mechanisms, dismissing important endogenous dynamics.

By varying the relative abundance A, it is possible to generate different levels of inequality. Panel (a) in Fig. 5 shows different limit distributions of housing ownership under various levels of relative abundance of the common asset. As shown in this plot, the distribution of housing becomes more skewed as the common asset becomes scarcer. Panel (b) offers a more general perspective by plotting the Gini coefficient as a function of A. Similarly, the average transaction price also has a negative relation with A, but it is insensitive at the extremes (panel c). Finally, panel (d) shows how the composition of transactions changes with the level of A.

These results provide a proof of concept of the model. Nevertheless, it may still not be entirely clear how it generates inequality endogenously. I elaborate on this issue in the next section.

3.2. Causes of inequality in the model

Allow me to divide the sources of inequality into three: (1) the interaction process, (2) sub-optimal transactions and (3) agent heterogeneity.[10] For clarity of exposition, let me assume a benchmark setup in which agents are homogeneous, there is no economic behavior, and there is no birth-death process. As I progress in my explanation, I will add these elements.

3.2.1. Interaction process

Suppose that, whenever two agents meet, one of them (chosen at random) transfers a random fraction of their asset to the other. This is a well-known stochastic process that gives rise to skewed distributions of the common asset, even when all agents start with the same endowment. In the 1980s, sociologist John Angle (Angle, 1986) termed this

[10] A fourth source is parameter levels and the way these are configured. However, this is complementary to the previous ones; it could be considered as part of parameter heterogeneity; and is conditional on source 1.

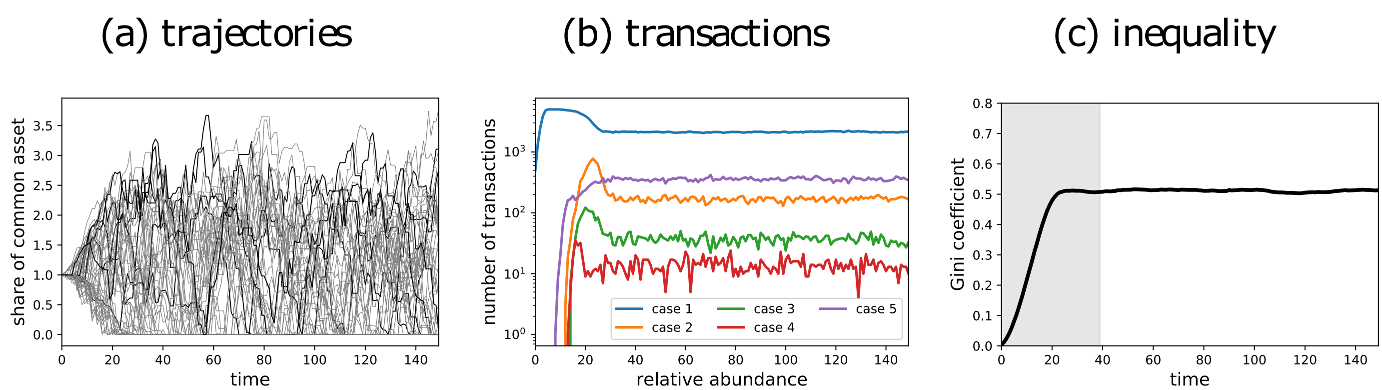

Fig. 4. Illustrative outputs.
In panel (a), black and grey lines represent the top-10 and the rest of the population respectively. These percentiles were calculated by averaging the agents' ownership of A across time. The shaded area in panel (c) denotes the transient of the simulation.

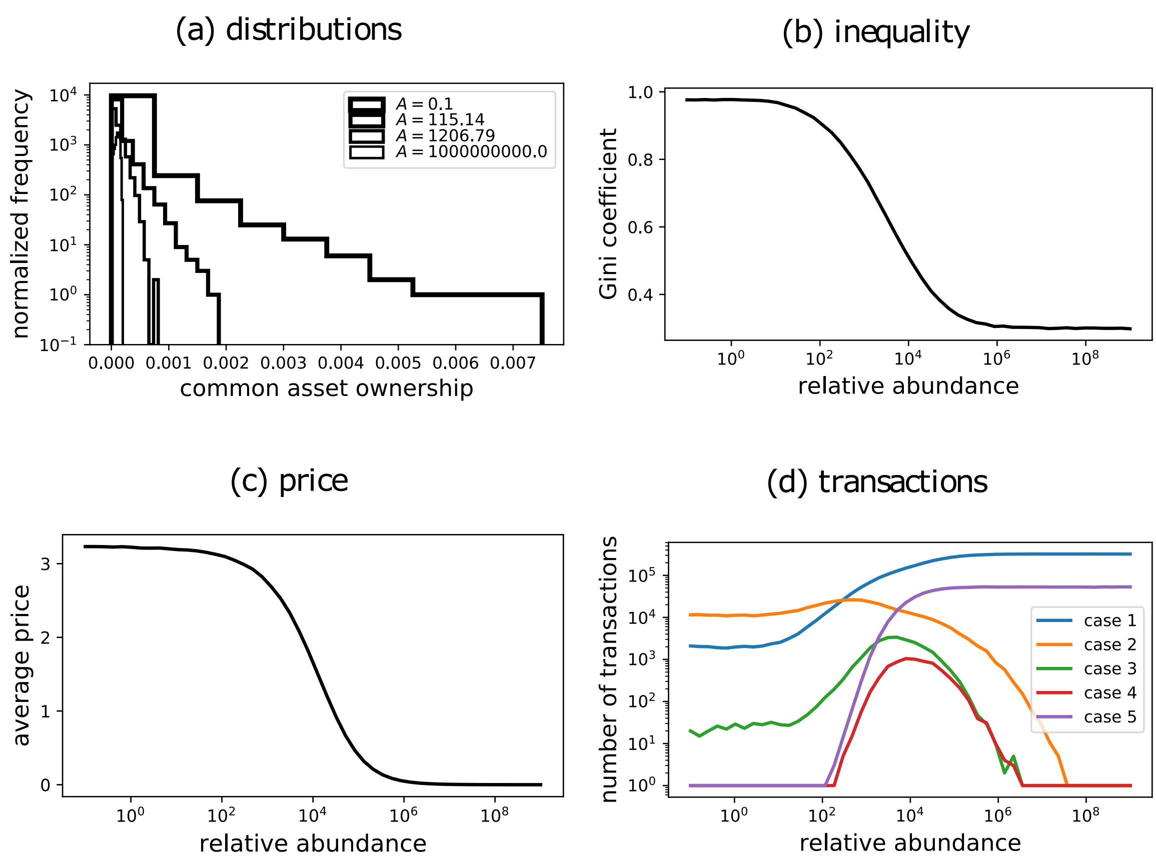

Fig. 5. HWI under different levels of relative abundance.
Panel (a) shows snapshots of the housing distribution during the last period of simulations with different levels of A. The rest of the panels were obtained by averaging the relevant statistic across the last 50 periods of each simulation (already in the steady state). The price in panel (c) corresponds to the average transaction price of each period, which is different from a price index. The former is the result of particular interactions in a given period, while the latter should reflect a market-level valuation. Section 4.3 develops a price index for the model.

model the one-parameter inequality process (OPIP).[11] More recently, econophysicists have employed it to analyse income distributions (Chakraborti & Chakrabarti, 2000).

It is important to mention that the OPIP, in spite of generating skewed distributions, has been largely dismissed by social scientists. In an extensive review of this type of models, Lux (2005) explains that the OPIP lacks credible social foundations to justify why agents would unwillingly transfer wealth to each other. The model presented in this paper generates similar dynamics to those from the OPIP, but with economic micro-foundations.

In order to generate inequality through the interaction process, incentives are necessary. However, if the populations is perfectly homogeneous, such incentives cannot exist because the transaction space is undefined between pairs of identical agents (unless $h \neq 1$). Nevertheless, this is only a mathematical result that holds for exact parameter values, and it is highly unlikely to exist in the real world. In fact, given that all parameters are continuous variables, the probability

[11] I establish a direct connection between the OPIP and this model in Appendix C.

of finding perfectly identical agents is zero. One can think of two ways to allow natural stochasticity in the system: (1) the aging/birth/death process, and (2) noisy parameters. The aging/birth/death process produces natural stochasticity because age discrepancies generate differences in the way agents discount future utility. This enables initial transactions that lead to HWI (because A_i produces heterogeneity endogenously). If this process is suppressed from the model, the only source of stochasticity left is the random values that the agents' parameters could take due to idiosyncratic or circumstantial reasons. For example, suppose that, every period, housing preferences are determined by

$$\beta_{i,t} = \eta_i + \varepsilon_{i,t}, \tag{11}$$

where η_i is the estimated value of the parameter and $\varepsilon_{t,\, i} \sim N(0, \sigma_i)$ is a noise term. On average, housing preferences remain as estimated. However, the noise perturbations enable transactions between otherwise identical agents.

Ideally, σ_i should capture parameter uncertainty. For example, $\sigma_{i,\, \alpha}$ could provide information about the range of values that the estimated consumption preference for individual i could take. Likewise, the distribution of α_i does not need to be restricted to a normal, but it could have any form that the investigator may find suitable. Fig. 6, illustrates different outcomes produced by the model without the aging/birth/death process, under different levels of σ_i.

It is important to clarify that, while the reader may attribute HWI to the stochasticity of the model (from the aging/birth/death process or from parameter stochasticity), stochasticity by itself is unable to generate these dynamics without an interaction protocol. Thus, one should interpret these results as the contribution of the market institution under the system's natural stochasticity. In the empirical application, I further decompose these dynamics by separating both sources of natural stochasticity.

3.2.2. Sub-optimal transactions

Now, let's consider a setup with homogeneous agents interacting in accordance to the model. Additionally, assume that there are no sources of natural stochasticity. Two results stem from this configuration:

Proposition 1. If infinitely-lived homogeneous agents transact only in the optimum, the steady-state distribution of the common asset is homogeneous.

Proposition 2. If infinitely-lived heterogeneous agents transact only in the optimum, the steady-state distribution of the common asset is fully determined by the agents' characteristics.

Appendices B.2 and B.3 provide the respective proofs of these propositions. Proposition 1 establishes that, even with heterogeneous initial endowments of the common asset, inequality vanishes under homogeneity and optimality. Proposition 2 establishes that, under a setting with optimal transactions, inequality can only be explained by agent heterogeneity, not by the interaction process (i.e. the market institution). Both results are quite profound because they suggest that assuming full rationality may limit the model's capacity to account for institutional factors. Thus, sub-optimal transactions allow inequality to emerge endogenously, even when agents are homogeneous.

3.2.3. Agent heterogeneity

Proposition 2 suggests that agent heterogeneity is important to explain HWI. In the presence of sub-optimal transactions, it contributes to the overall level of inequality by interacting with the behavioral component that determines equilibria. The influence of heterogeneity varies depending on the attribute that is being varied, and on the fraction of the population that is assumed heterogeneous. Allow me to illustrate this through wages.

Fig. 7 presents three simulation outputs (one per column) with different levels of wage heterogeneity. Here, I have assigned a percentage of agents with a randomly-allocated wage (using a uniform distribution in (0,1)). The top panel in column (a) shows some reallocations of the common asset at the beginning of the simulation. However, eventually, all the agents end up with the same share. The bottom panel of the same column suggests that this is because optimal transactions become the only ones in the market. Thus, the outcome is the homogeneous distribution predicted by Proposition 1.

Panel (b) shows more reallocations and transaction cases. Here, the economy reaches some level of inequality and a steady proportion of sub-optimal transactions. Finally, panel (c) shows the outcome for full wage heterogeneity. Here, all transaction cases have occurred at some point in the simulation, and inequality seems more pronounced as the trajectories of the top-20 (thin lines) are clearly separated from those of the rest (thick lines).

Finally, I would like to show that the level and type of heterogeneity causes inequality in non-trivial ways. For this, I use three types heterogeneity: (1) wages, (2) discounting, and (4) housing preferences. Wages are captured by w_i and affect the budget constraint directly. Discounting is represented through γ_i and influences the perception of future utility; its effect is similar to the one that age differences would have. And housing preferences are represented by β_i.

For each type of heterogeneity, I maintain all other parameters homogeneous, set homogeneous initial endowments of the common asset, and block the aging/birth/death process. The exercise consists of activating only one type of heterogeneity and varyingthe proportion of the population that is heterogeneous. For the heterogeneous agents, the relevant parameter is assigned at random (note that the four parameters

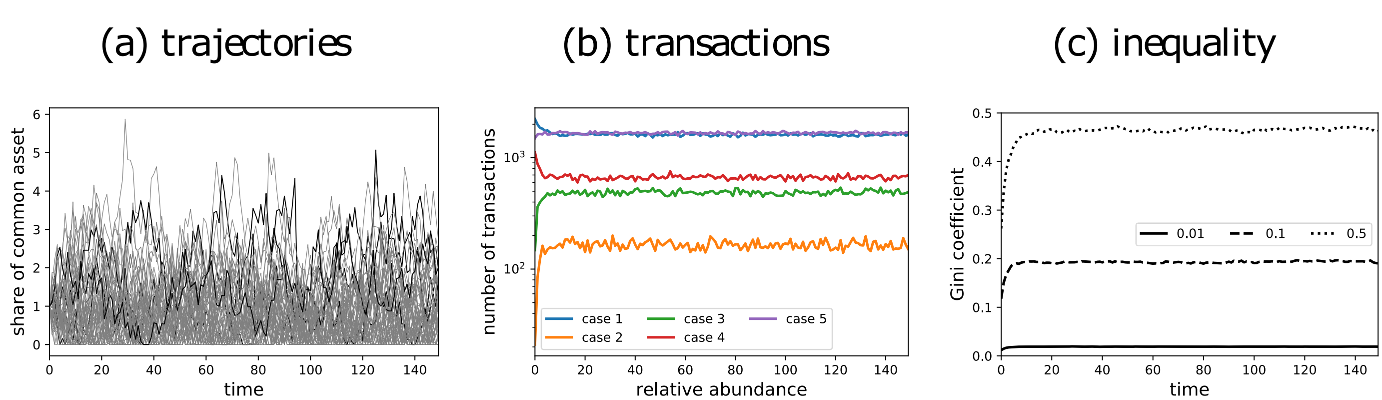

Fig. 6. Dynamics through the interaction process.
Model dynamics in a homogeneous, non-aging, infinitely lived population under different levels of parameter stochasticity. Panels (a) and (b) correspond to a single illustrative run. Panel (c) shows the aggregate dynamics of three simulations with different levels of parameter stochasticity. In panel (a), black and grey lines represent the top-10 and the rest of the population respectively. These percentiles were calculated by averaging the agents' ownership of A across time.

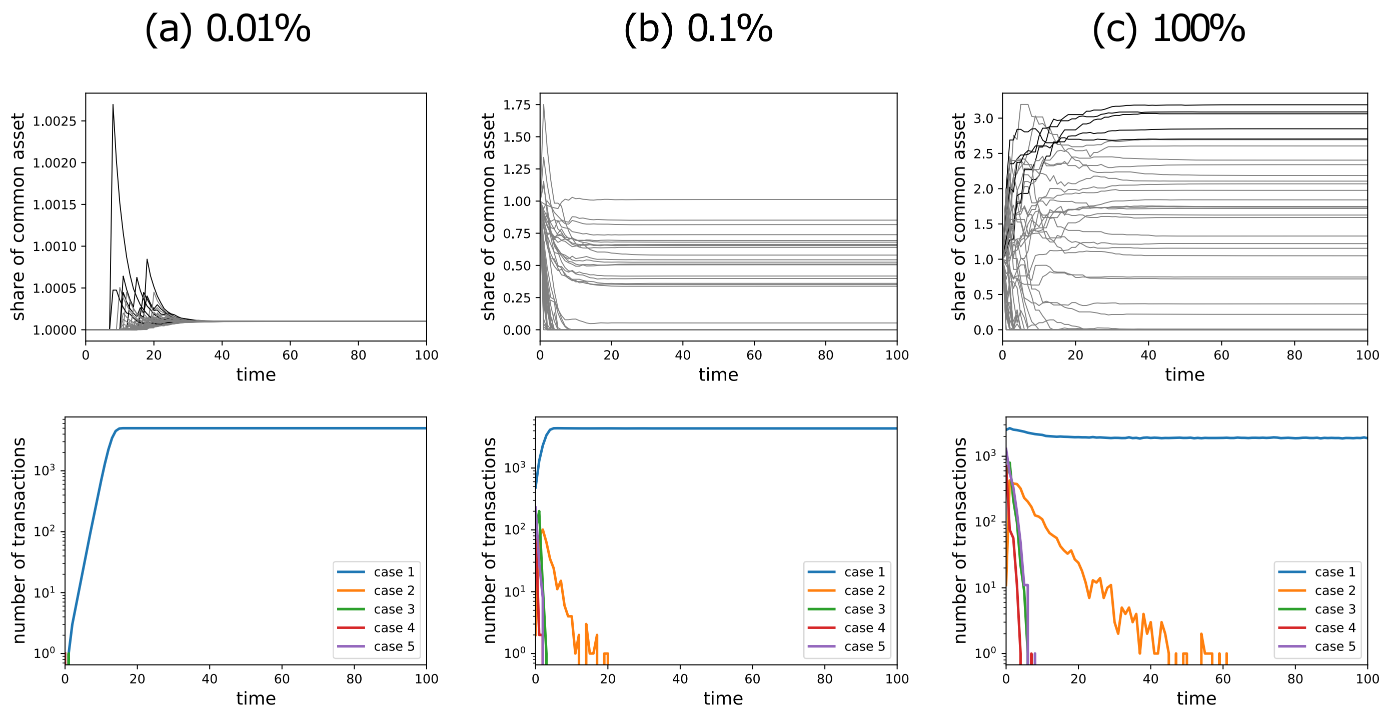

Fig. 7. Dynamics under different levels of wage heterogeneity.
In the top panels, black and grey lines represent the top-10 and the rest of the population respectively. These percentiles were calculated by averaging the agents' ownership of *A* across time.

live in (0,1)). A simulation is run, and the statistic of interest is captured for the steady state. This procedure is performed multiple times to obtain an average statistic for different randomizations of the same type. Finally, all the previous steps are performed for each type at different heterogeneous population proportions.

The left panel in Fig. 8 shows the Gini coefficient emerging for different types and levels of heterogeneity. Clearly, more heterogeneity produces more inequality. However, the rate at which inequality increases differs by type.

The right panel in Fig. 8 shows the number of transaction cases in the steady state of a fully heterogeneous population for each type. The plot suggests important differences in the composition of transactions. For instance, in this example, the discounting type allows all transaction cases, while the other types do not exhibit cases 3 and 4.

To summarize, the model is able to disentangle different sources of HWI. All of them are important and interact in non-trivial ways. This is why an agent-computing approach is suitable for the empirical analysis of HWI. In the next section, I present an application with the same scale and heterogeneity as the British housing market.

4. Application

This section demonstrates the applicability of the model using empirical data. First, I provide a rough estimate of how much the market institution (the interaction process) accounts for HWI in each region of Great Britain. Then, I perform a series of policy experiments about interventions that are commonly discussed in the media and academia. Since, at this stage of development, the model lacks elements of housing markets such as space, explicit housing units, financial instruments, and a rental market, the policy experiments should not be interpreted as evaluations or prescriptions. Instead, they should be read as counterfactual simulations providing an intuition of how certain interventions could play out in the context in which the data is embedded.

4.1. The data

The UK Office of National Statistics (ONS) conducts the Wealth and Assets Survey (WAS) on a biannual basis. This survey measures household and individual wealth across multiple dimension, for example, value of properties, financial assets, labor income, capital income, etc. In addition, it allows the direct inference of differentiated tax rates by reporting income measures before and after tax. I use data on

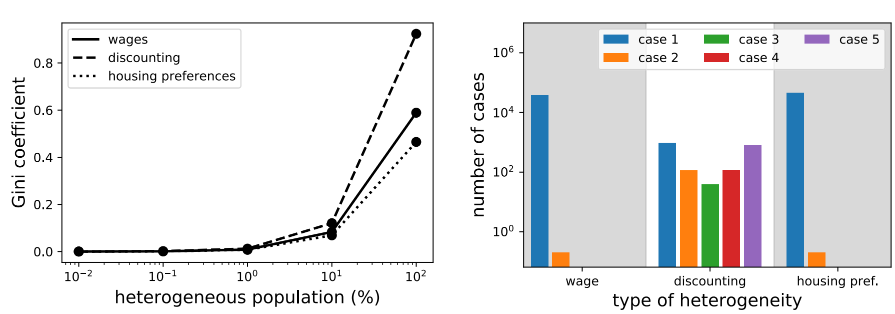

Fig. 8. Inequality generated by different parameters.
Left: average Gini coefficient per type and level of heterogeneity. Right: average number of transaction type under each parameter with 100% heterogeneity.

the fifth wave of the WAS (from July 2014 to June 2016) in order to instantiate a synthetic population of households with the characteristics of the WAS sample. That is, each one of the ∼25 million households in Great Britain is simulated and individually calibrated.[12]

The data covers 11 regions of Great Britain: North East (NE), North West (NW), Yorkshire and the Humber (YH), East Midlands (EM), West Midlands (WM), East of England (EE), London (LN), South East (SE), South West (SW), Wales (WL), and Scotland (SC). I calibrate the model to each of these regions, assuming that they are independent from each other, which significantly reduces the computational burden of such a large population of agents and of their interactions. Table 2 summarizes the WAS data across regions.

The consumption preference parameter α, was estimated according to the identity

$$\alpha_i = \frac{c_i}{\tau_i w_i \ell_i + c_i}, \tag{12}$$

derived from the first-order conditions of the utility maximization problem in Eq. (2).

The ONS provides a separate dataset with the annual discount rates used through each interview stage of the WAS.[13] From these data, the average discount factor in the 2014–2016 period can be estimated as $\gamma = 0.965$.

Regarding the survival probability δ, the ONS produces the National Life Tables, where estimates are provided conditional on age and gender. Because the model allows a highly dissagregate calibration, I employ these data. Therefore, the calibration accounts for gender differences in life expectancy for different age groups.[14]

Housing preferences β_i were inferred through the seller's indifference curve, defined in Eq. (7). Using data (from the WAS) on the number of properties owned by the interviewees, I assume a hypothetical situation in which the agents could sell the entirety of their asset for the values reported in the data. The fact that this situation has not realized implies that the utility surplus from doing it is null; hence the relevance of the indifference curve. Thus, by solving Eq. (7) for β one obtains the identity

$$\beta = \frac{hp}{hp + \mu_L A(\tau w + zB + s)}, \tag{13}$$

where A and p are the number and total value of the properties owned by the agent (as declared in the WAS). All the RHS elements in Eq. (13) are obtained from the WAS. This parameter, however, can only be defined among those who have reported owning properties and their values. For those without this information or with $\beta = 0$, I assign the average value of β in the corresponding region.

4.2. Calibration

Table 3 presents the model variables with their sources and calibrated values. As a benchmark, I set the property sales tax to zero. The rest of the parameters are directly imputed from the WAS micro-data.

In order to introduce natural stochasticity in the model parameters, I employ the formulation from Eq. (11) that adds a noise term distributed according to a normal with mean zero and standard deviation σ. Every period, a random value for the noise term is drawn form this distribution and added to the parameter estimate. Ideally, the random draw should be done on a truncated normal, in order to comply with the parameter's boundaries. However, this is a considerable computational burden for the current implementation of ∼25 million agents. Therefore, I use a non-truncated normal, and realize the noise addition only if the draw falls within the boundaries of the parameter. In other words, the emergent inequality from parameter stochasticity may be lower than the one under a truncated normal because the parameters are perturbed less frequently. Nevertheless, I have found through small-scale simulations that, for this application, the difference between using a truncated normal and my heuristic is negligible.

In principle, σ should capture the uncertainty of the parameter in question. Since the model runs at a one-to-one scale, no such information is contained in the data (because each agent is a data point). Therefore, the second-best is to compute σ for a sample of agents that are ‘similar’. For this, I pool agents from the same region and from the same age group, and compute the variance of their parameters. This means that, at a given age, a certain agent in a region may experience a perturbation in their parameters for idiosyncratic or circumstantial reasons (but, on average, their parameter is the one coming from the data). Not all age groups have enough agents to compute σ, especially in the extremes of the age distribution, so these perturbations do not happen in every step of the agents' lives.

Next, I present the Gini coefficient of housing ownership, directly calculated from the WAS for the different regions of Great Britain. The left panel in Fig. 9 shows that the North East region has the highest levels of inequality in terms of property ownership. The lowest corresponds to the East of England, while London seats in the middle with a Gini of approximately 0.8.

By manipulating the relative abundance parameter A, it is possible to vary the model's (endogenous) Gini coefficient. I show these variations in the right panel of Fig. 9. Interestingly, the different ranges of the model's Gini coefficients are compatible with the empirical data. That is, in all the regions, it is possible to find a unique A such that the model generates the observed level of inequality. Therefore, calibrating the relative abundance consists of performing a greedy search until the difference between the empirical and the simulated Gini coefficient is minimal.

4.3. External validation

Before presenting the main result, I would like to argue for the validity of the model. While external validity may adopt different meanings depending on the discipline and method, here I define it as the ability of the model to reproduce a stylized fact that (1) is measured through an independent data source (the model parameters do not come from such data) and (2) is not a target metric to fit the model. The stylized fact that I use for validation is price heterogeneity across regions. This information can be obtained from the ONS House Price Index (HPI), which has been estimated for each British region, and is independent of the WAS. In contrast with the WAS –which relies on self-reported values–, the HPI is built from administrative data held by Her Majesty Land Registry, and Registers of Scotland. My aim is to evaluate whether a measurement of price in the model shows a strong correlation with the HPI across the British regions. First, by computing different types of correlations, I evaluate validity and robustness. Second, by measuring the correlation between the HPI and other variables from WAS, I check if the validity is not trivially driven by other factors in the data; i.e. I evaluate if the model generates new information.

In order to construct a price index for the model, it is necessary to obtain valuations of the shares held by the agents. The procedure consists of creating an ‘average’ agent in each region; a synthetic household with the average characteristics of the region's population. This agent approaches each household with the intent of purchasing the entirety of their common asset. This is equivalent to computing the equilibrium price in Eq. (2).3 by using $q = A_j$, and assuming that i is the average agent and j is the seller. The price index is the average of these elicited prices.

Table 4 reports the outcome of this exercise. First, it shows that both the model price index and the data on wages have a strong linear

[12] While the WAS consists of a sample of British households, it is possible to use the variables' weights to generate a synthetic population. Such weights are already corrected for sampling biases.

[13] These data can be downloaded from the ONS website by searching for *annuity rates and discount factors, July 2014 to June 2016*.

[14] The National Life Tables can be obtained from the ONS website.

Table 2
Empirical data from the Wealth and Assets Survey by region.

Region	N	w (£)	B (£)	s (£)	$1-\tau$ (%)	$1-z$ (%)	age	α	β
NE	1,002,792	23,232.78	137.28	406.21	18.12	0.02	45.34	0.25	1.08
		(16,916.82)	(961.43)	(1472.12)	(11.8)	(0.77)	(11.71)	(0.09)	(0.94)
NW	2,876,133	23,521.34	643.72	337.24	18.64	0.2	46.22	0.25	1.0
		(18,015.44)	(9798.05)	(1003.8)	(12.17)	(2.59)	(11.95)	(0.09)	(0.97)
YH	2,129,157	24,012.08	337.11	407.72	18.67	0.15	45.81	0.26	0.98
		(19,259.62)	(4801.83)	(1193.74)	(12.03)	(1.99)	(11.81)	(0.09)	(1.0)
EM	1,932,563	24,480.16	332.27	362.88	19.19	0.18	45.31	0.26	0.97
		(21,035.87)	(2637.7)	(1108.29)	(11.87)	(2.14)	(12.48)	(0.09)	(1.84)
WM	2,267,401	23,366.94	119.32	389.73	18.9	0.03	47.19	0.25	0.86
		(21,685.0)	(693.6)	(1158.83)	(11.83)	(0.76)	(11.13)	(0.09)	(0.64)
EE	2,458,136	27,784.59	554.0	338.96	19.71	0.28	45.94	0.26	0.8
		(26,245.52)	(5154.71)	(1193.97)	(12.08)	(2.8)	(11.35)	(0.1)	(0.85)
LN	3,344,438	37,324.02	1755.18	401.09	22.19	0.63	47.13	0.27	0.69
		(62,083.05)	(19,203.72)	(1465.02)	(12.82)	(4.47)	(11.56)	(0.1)	(1.49)
SE	3,704,864	31,730.53	1585.28	296.12	20.84	0.34	46.96	0.27	0.7
		(34,408.47)	(84,753.85)	(984.27)	(12.69)	(2.95)	(11.16)	(0.1)	(0.6)
SW	2,219,704	24,768.37	448.75	334.35	19.1	0.19	46.84	0.26	0.75
		(21,400.81)	(5099.8)	(980.23)	(12.08)	(2.22)	(11.82)	(0.1)	(0.55)
WL	1,200,342	22,500.56	130.01	339.11	18.64	0.03	46.14	0.25	1.01
		(16,377.72)	(715.01)	(1106.52)	(11.69)	(0.59)	(11.53)	(0.08)	(1.81)
SC	2,188,172	26,256.41	220.27	397.2	19.38	0.06	46.23	0.25	1.07
		(24,757.72)	(2378.31)	(1311.24)	(11.51)	(1.32)	(11.29)	(0.09)	(1.39)
GB	25,323,702	27,345.49	724.97	360.15	19.69	0.24	46.42	0.26	0.87
		(32,129.82)	(33,438.39)	(1178.33)	(12.22)	(2.61)	(11.6)	(0.1)	(1.14)

Average value of each variable per region and its corresponding standard deviation in parenthesis.

Table 3
Model calibration.

Parameter	Concept	Source	Estimation
A	Total common asset	Via simulation	To calibrate
N	Number of agents	Population size	~25 million
α_i	Consumption preferences	ONS & model	Heterogeneous
β_i	Housing preferences	ONS & model	Heterogeneous
γ	Discount rate	ONS	0.965
δ	Survival probability	ONS	Heterogeneous
w_i	Wage	WAS	Heterogeneous
s_i	Government transfers	ONS	Heterogeneous
B_i	Non-labor income	ONS	Heterogeneous
$1-\tau_i$	Labor income tax	ONS	Heterogeneous
$1-z_i$	Non-labor income tax	ONS	Heterogeneous
$1-h_i$	Sales tax on common asset	Benchmark	0.000
c_i	Consumption	Endogenous	NA
$1-\ell_i$	Labor	Endogenous	NA
A_i	Common asset ownership	Endogenous	NA

Non-labor income is obtained from the WAS variable: "household gross annual income from investments".

Table 4
Model validation through the House Price Index.

Variable	Pearson r	Spearman r	Kendall τ
Model price	0.928	0.973	0.891
	(3.81E-5)	(5.14E-7)	(1.37E-5)
Wage	0.947	0.736	0.636
	(9.92E-6)	(9.76E-3)	(5.71E-3)
Non-labor income	0.886	0.664	0.527
	(2.79E-4)	(2.60E-2)	(2.64E-2)
Government transfers	−0.12	−0.382	−0.309
	(7.26E-1)	(2.47E-1)	(2.18E-1)
Gini coefficient	−0.252	−0.482	−0.418
	(4.55E-1)	(1.33E-1)	(8.66E-2)

The correlations were computed between the variable of interest and the ONS House Price Index during the sampling period. The unit of observations are Great Britain's regions. The numbers in parentheses correspond to the p-values of the correlations.

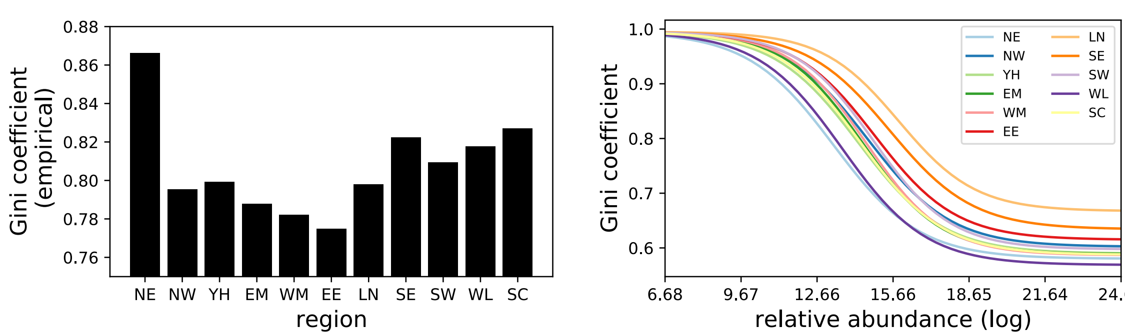

Fig. 9. Empirical and simulated HWI.
Left panel: empirical Gini coefficients. Right panel: simulated Gini coefficient as a function of the relative abundance.

correlation with the HPI. However, when one moves to non-linear correlations, the model improves its performance while the data on wages drops substantially. Other income-related variables that, arguably, can drive house prices show even weaker correlations to the HPI. Finally, because the model's free parameter is calibrated against the Gini coefficient, it is important to check whether the Gini also shows a strong correlation with the HPI. The table shows that this is not the case. Therefore, the model is not only able to generate endogenous

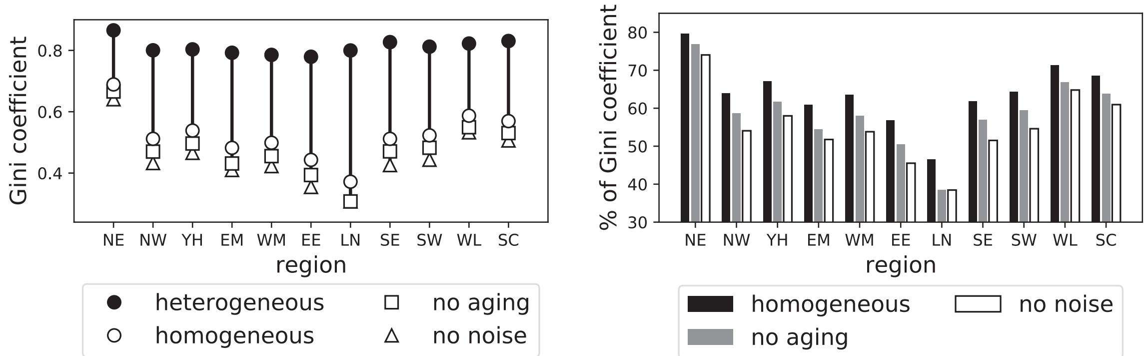

Fig. 10. Housing wealth inequality and the market institution.
Left: Gini coefficients estimated for control (solid markers) and treatment groups. Right: Gini coefficient of treatment groups as a fraction of the one empirical one.

HWI, but also to reproduce the independent HPI better than what could be done through back-of-the-envelope calculations using income data. This demonstrates the external validity for the model.

4.4. On the causal inference of the market institution

As I have previously argued, computational models of this sort are excellent tools for the causal inference of institutional factors. For instance, one can entirely remove a tax or a regulation in a synthetic population (the treatment group), while leaving it in another one with identical agents (the control group). When trying to apply this logic to the market institution, however, one quickly realizes that this is an ill-posed problem.

Suppose one removes the market from the model; what would take its place? Any substitute that involves voluntary transactions certainly has a market element. Furthermore, one could think of endless alternative non-market allocations of A, for example, a uniform one, one that correlates A_i with w_i in a negative way, or even a random allocation. Which one should one choose? Which one provides a suitable benchmark? There does not seem to be an answer that eludes arbitrariness, so this causal inference strategy does not seem to work, even under the benefits of a computational framework.

What alternative is there to infer market effects? In order to answer this, one needs to go back to the sources of HWI discussed in Section 3.2; more specifically: heterogeneity. The interaction process and sub-optimal transactions are integral parts of the market institution, as modeled in this paper. Thus, by removing the source of heterogeneity, the remaining sources of HWI correspond to the market institution.

The strategy, thus, consists of creating a treatment group in which all agents are homogeneous (they have the average parameters of the region). Such group is let to interact in a simulation with the level of A that was originally estimated for its region (its level of relative scarcity). The inequality that emerges from this simulation is attributed to the interaction process and its sub-optimal transactions, i.e. to the market institution. Thus, the difference between the control and the treatment groups is the market effect. Further refinements of this effect are presented by removing some of the market's natural stochasticity.

4.5. Market effects

The calibrated model that matches the empirical Gini provides a control group through its synthetic population of artificial agents. As explained above, in order to estimate the effect of the market institution, I design another type of simulations in which heterogeneity is removed from the agents' parameters. The agent population of this simulation provides the treatment group. Moreover, I setup three different treatment groups by specifying three counter-factual simulations in order to flesh out the contributions of the two sources of natural stochasticity in the model: (*i*) one with homogeneous agents and both sources of natural stochasticity, (*ii*) another with homogeneous agents but no aging/birth/death process, and (*iii*) a third one with homogeneous agents but no parameter stochasticity. The treatments are run in the same fashion as the control group: providing identical initial endowments and letting the agents interact until the steady state is reached. The first treatment represents the inequality that emerges from the market institution only, because all heterogeneity has been removed from the population, with exception of the natural stochasticity. The second treatment captures the inequality that would emerge from the market institution in a world in which the aging/birth/death process does not produce heterogeneous expectations. The third treatment measures the inequality that the market would produce in a population without idiosyncratic or circumstantial perturbations to the agents' parameters. Thus, the first treatment is the main result, while the other two are refinements to understand how the different sources of stochasticity enable the endogenous formation of HWI.

Fig. 10 presents the results of this exercise. The left panel shows the estimated levels of the Gini coefficient, while the right one shows the Ginis of the treatment groups as a fraction of the one from the control group. Several fascinating outcomes can be identified. First, there is substantial heterogeneity in the effect that the market has on HWI across Great Britain. For example, the market seems to be responsible for 40 to 50% of the Gini in London, while it accounts for more than 70% in the North East region. This means that agent heterogeneity plays a more predominant role in explaining HWI in London than in other regions (presumably, income inequality). Second, parameter stochasticity is a stronger enabler than the aging/birth/death process (the triangular markers are always the lowest in the left panel). Third, removing the aging/birth/death process does not necessarily mean lower HWI. Note that the squared markers in the left panel can be higher or lower than the circle ones.

Overall, according to this model, it can be said that the market is responsible for one to two thirds of the Gini coefficient measuring HWI across Great Britain. The model provides a causal account for HWI, and the experimental design allows identifying counterfactuals in which agent heterogeneity (the traditional explanatory variables of HWI) is effectively removed. While this is still a rough estimate (because more realistic assumptions need to be incorporated into the model), it provides a novel approach to measure the market origins of HWI. Thus, this is the main result of the paper.

Now that a causal account for HWI has been established, the reader may be interested in the potential outcomes from policy interventions. Agent-computing models are particularly well enabled to perform policy experiments. Therefore, I present results on three commonly discussed types of interventions: an expansion of the housing stock, a sales tax, and an inheritance tax. The results should be interpreted on a more qualitative basis (focusing on their signs rather than on specific numbers). Should the reader wish to obtain more precise estimates, further details on the policy channels and, in particular, on how taxes

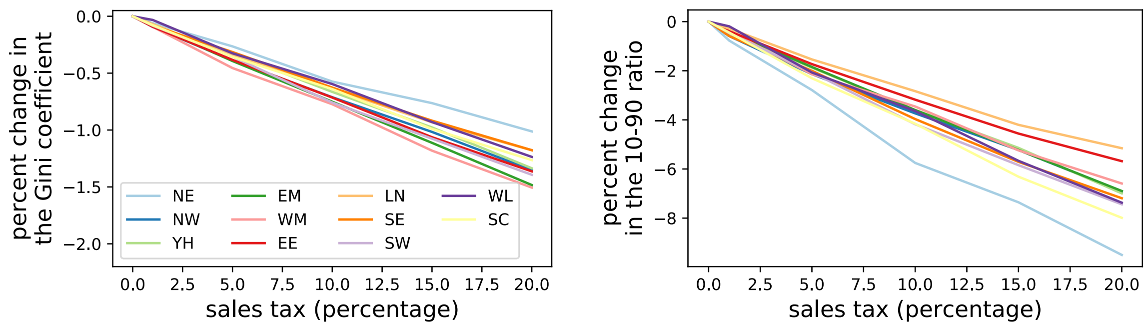

Fig. 11. Housing stock growth and HWI.
Left panel: changes in the Gini coefficient. Right panel: changes in the ratio of total housing concentrated in the top 10% against the bottom 90%.

are redistributed should be specified in the model. Needless to say, no statistical approach is except of this limitation so, while still some work needs to be done in this regard, this model provides a transparent way to incorporate the nuances of certain institutional settings, something impossible when analyzing aggregate statistical relations.

5. Policy experiments

I focus on three policies that are commonly mentioned in British public debates: (*i*) an expansion of the housing stock, (*ii*) a sales tax, and (*iii*) an inheritance tax. An expansion of the housing stock is the most commonly argued solution in political discourse. Thus, it is interesting to examine the hypothetical situations in which the stock of the common asset A increases. Of course, the model does not consider the political economy dimension of how the government incentivices developers, an interesting aspect that I leave for future research. Policies 2 and 3 have a more direct interpretation in the model since h accounts for the tax extracted from a sale, and an inheritance redistribution scheme can be easily implemented in a stylized fashion. Below, I provide further details on the design of each experiment and present the main findings.

5.1. Expansion of the housing stock

The control group in this experiment is the artificial population where A has been calibrated to match the empirical Gini. The treatment group, on the other hand, is the same artificial population, but subjected to an augmented housing stock A'. The treatment has to be run independently of the control because there are no explicit economic mechanisms to grow A during a simulation. That is, while it is technically possible to augment A *during a simulation*, some assumptions regarding the initial ownership of the new common asset would be necessary, for example, to incorporate a 'developer' agent. Since such specificities lie beyond the scope of this paper, I take a similar approach to the one used in Section 4.5.

I explore the outcomes from increasing the common asset in 1, 5, 10, 15 and 20% in each region. Note that most of these growth rates are unrealistic policies. For instance, according to the Department of Communities and Local Government (DCLG, 2016), the number of dwellings in England alone was approximately 23 million. On average, this housing stock grew 0.77% each year during the sampling period; less than the most conservative scenario proposed here. Yet, it is useful to study these hypothetical scenarios since the effectiveness of expanding the housing stock has been questioned by housing scholars (Gallent, Durrant, & May, 2017; Ryan-Collins, 2019).

In addition to presenting the effects on the Gini coefficient, I report changes in the rate of accumulated housing ownership between the top 10% and the bottom 90%. This metric is useful to understand the disparities between the most wealthy and the rest of the population.

Fig. 11 shows the result of the experiments. While expanding A seems to have a negative impact on the Gini coefficient, the overall effect is quite moderate, not exceeding a 2% decrease in any of the regions. In addition, the effectiveness of this policy instrument seems to vary across regions, as suggested by the differences in slopes between the lines potted in the left panel. Interestingly, the effects are nonlinear, as some of the curves in the plot intersect at different increment levels of A. The right panel shows a decrease in the 10–90 ratio with the expansion of A. Something to highlight is that, in the North East, a change in A seems to be less effective on the overall housing distribution, however, this same region is the most sensitive in the 10–90 ratio.

5.2. Sales tax

The experimental design with regard to taxes is different from the ones previously performed. Here, the artificial population is directly intervened once it has reached the steady steady. The intervention is in the form of an instantaneous increase to the tax rate. Then, the simulation keeps running until the next steady state is reached. Finally, the analyses report the difference between the relevant variable in the two steady states: before and after the intervention. This design is preferred over the control-treatment one because it is closer to the way in which real-world dynamics play out (with agents adapting to the intervention).[15] Of course, some assumptions are need with regard to how taxes are collected and how they are redistributed.

The sales tax rate has already been specified in the seller's utility surplus (Eq. (5)) through $1 - h_i$. By default, this tax is set to zero,[16] so I explore rates of 10, 20, 30, 40 and 50%. Here, I assume that redistribution takes place through the government transfers parameter s_i. The implementation is straightforward. At the beginning of every period, the government distributes all the sales tax that it collected in the previous step. This redistribution increases parameter s_i. Next, all transactions take place, so the government refills the tax pot by extracting $1 - h_i$ from each sale. Finally, all agents' government transfers are set back to their baseline empirical levels. This implementation assumes another time inconsistency: the agents behave as if the increased s_i will remain the same in the future (which, on average, holds in the steady-state).

Fig. 12 shows the results of the experiment. In general a sales tax would increase in inequality. Sales taxes, as redistributive policies, are usually considered regressive because they adversely and disproportionately impact the purchasing power of the poorest individuals. It is very interesting to see, however, that this reflects in the

[15] This exercise only looks at inequality as a consequence of taxation policies. Thus, it does not intend to assert that certain sales tax levels are generally better or worse. This would require taking into account other outcome variables and to perform wealth analysis.

[16] The UK does not have a sales tax on property, only a capital gains tax. With an infinitely divisible common asset, it is not possible to implement capital gains, since these are usually attached to specific properties.

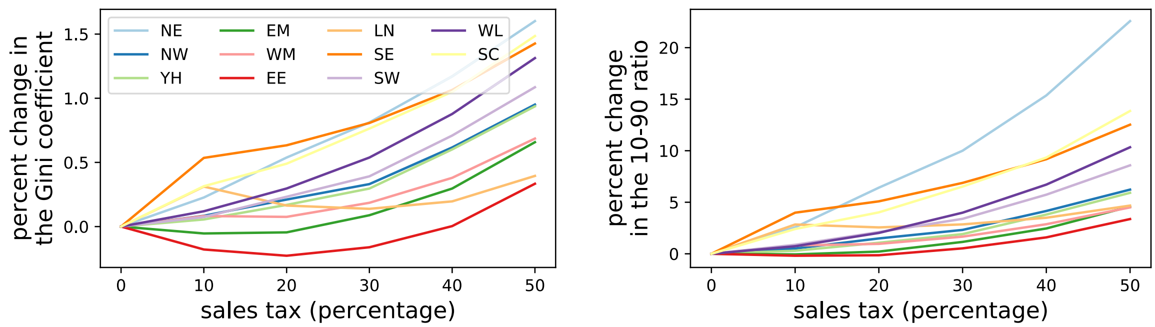

Fig. 12. Sales tax and HWI.
Left panel: changes in the Gini coefficient. Right panel: changes in the ratio of total housing concentrated in the top 10% against the bottom 90%.

outcome of the experiment in terms of housing consumption. In some cases, like in East England and London, certain tax levels reduce the Gini coefficient but, eventually, inequality raises again. When looking at the 10–90 ratio, there are also non-linearities. However, the case of the North East region stands out as a case there the top 10 would disproportionately benefit from such policy. These results suggest that the effect of a sales tax is not linear and not monotonic. They emphasize the importance of modeling institutions explicitly, as these non-linearities can produce unintended outcomes.

5.3. Inheritance tax

The last experiment looks at the effect of implementing an inheritance tax, a policy instrument that has been highly debated in recent years, not only to curve HWI, but wealth inequality in general.[17] Inheritance taxes can be implemented in this model due to the aging/birth/death process. Whenever an agent i dies, a fraction k_i of their common asset is taken by the government, and the remaining $(1 - k_i)A_i$ is inherited by their successor. By the end of each period, the government redistributes the total common asset collected by directly allocating an equal fraction to each agent. This, of course, assumes that the tax is collected and redistributed in specie.

Fig. 13 presents the results. As expected, there is a negative effect in the Gini coefficient. The effects are substantially larger in this case, however, one should be careful in not interpreting these results in terms of more or less effectiveness. The experiment assumes collection and redistribution in specie, so it is entirely possible that, when introducing more realistic collection and redistribution mechanisms, these magnitudes change. Nevertheless, it is interesting to confirm that the intuition that inheritance tax reduces the Gini. In the right panel, one can see that the response of the 10–90 ratio is much more aggressive than the one from the Gini. This is interesting because it seems that the different housing markets may have different sensitivities to the types of interventions. In the next section, I explore this.

5.4. Comparing interventions

While statistics such as the Gini coefficient or population ratios provide insights about the changes in inequality, it is still useful to look at the entire housing ownership distributions and how it changes under the different policy interventions from the previous experiments. This is important because a policy may not necessarily target an aggregate index, but rather certain sub-populations such as the most disadvantaged ones (e.g. social housing). By comparing the histograms of an example region (East of England), Fig. 14 shows the response of the distribution to the three policy interventions across their different levels. The result is fascinating.

Although an expansion in the housing stock reduces the Gini, this is not because the upper tail of the distribution has shrunk, but because the entire distribution has shifted to the right. Thus, with more common asset in the market, the most disadvantaged groups can buy more of it, but so do the wealthiest ones. Overall, the reduction in inequality comes from the increase of housing in the lower tail; not from a progressive redistribution.

The distribution corresponding to the sales tax shows an expansion in the upper tail and a contraction in the lower one. The already-wealthy agents become wealthier while the poorest become poorer thanks to this policy, confirming the –commonly argued– regressive nature of these taxes as redistributive tools. In contrast, the inheritance tax generates a contraction of the upper tail. This change is more in line with redistributive principles that intend to limit the prevalence of the 'super wealthy'. Therefore, when suggesting policy prescriptions, it is important to not only look at aggregate statistics, but also at the disaggregate composition of the variables of interest. Furthermore, policy regimes that intend to mix different types of policies should be worry of non-trivial interactions and non-linearities.. The histograms of all the regions across the three policy experiments can be found in Appendix F.

Finally, it is interesting to compare the sensitivity rankings that the different regions have to each policy. That is, for a given policy experiment, I rank each region according to the average absolute change in the Gini coefficient. Then, I compare how these ranks change from policy to policy. Rank changes are interesting because they reveal something about the differentiated effects of the three interventions and, hence, highlight the importance of policy design that is tailored to specific contexts.

Fig. 15 shows the sensitivity rankings of the three interventions, being 1 the most sensitive region and 11 the least one. Sensitivity to an intervention is measured as the average change in the Gini coefficient (across the different levels of the same intervention) per 1% change in the intervened variable. Overall, it is interesting to see that the rankings are not static across the three experiments. For instance, the West Midlands are the most sensitive to a housing stock expansion, but rank at the bottom of taxation policies. London, on the other hand, is not very sensitive in the first two experiments, but then becomes the fourth most sensitive when it comes to inheritance taxes. This exercise highlights the importance of understanding differentiated regional effects because, when it comes to establishing national-level policies, one would like to choose the rules of the game that maximize overall effectiveness.

6. Conclusions

While housing wealth inequality has become central in British public debates, we still lack robust evidence of which policies could provide reliable solutions. In part, this is due to difficulties in measuring market effects in a system for which we do not have natural experiments. This paper introduces an economic computational model that

[17] Appleyard and Rowlingson (2010) study the linkages between inheritance, housing wealth and inequality.

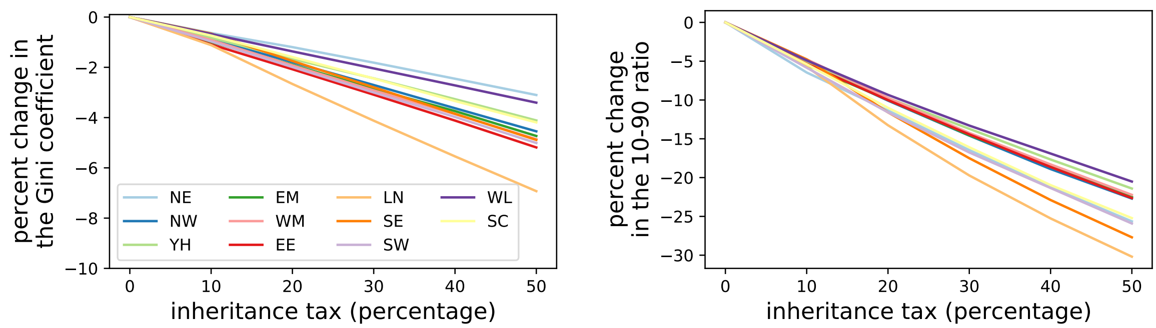

Fig. 13. Inheritance tax and HWI.
Left panel: changes in the Gini coefficient. Right panel: changes in the ratio of total housing concentrated in the top 10% against the bottom 90%.

Fig. 14. Distribution of housing ownership under policy interventions.
The gradient of the lines decrease with the size of the intervention. The size of the intervention is in percentage.

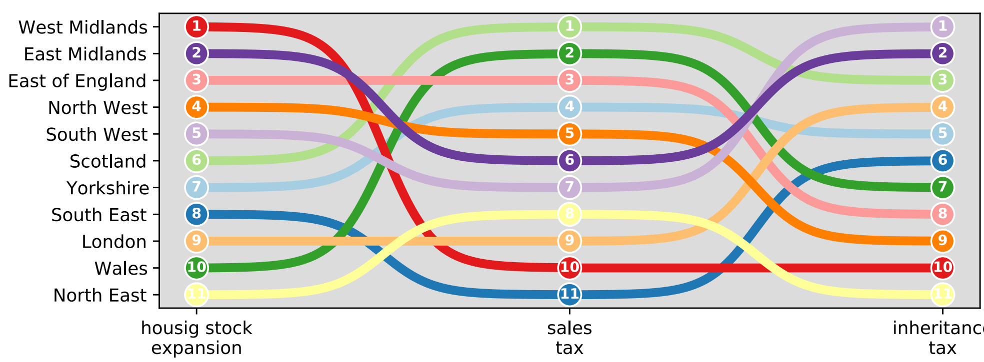

Fig. 15. Sensitivity rankings.
Sensitivity to an intervention is measured as the average absolute change in the Gini coefficient (across the different levels of the intervention) per 1% change in the intervened variable.

facilitates the measurement of market effects on HWI. The model is able to control for all observed exogenous sources of variation, and simulates the entire household population to generate HWI endogenously. I provide rough estimates on how much HWI is attributable to the market, i.e. to the *tatonement* process through which British households buy and sell property. I find that, across 11 regions, the market causes between one and two thirds of the Gini coefficients.

I perform policy experiments using intervention instruments that have been considered in public debates: an expansion of the housing stock, a sales tax, and an inheritance tax. I find that each of these policies yield different levels of effectiveness and their impact varies significantly across the different regions.

In spite of these results, the model is still highly stylized as it assumes that housing is an infinitely divisible good, that there are no financial instruments, that there is no spatial structure within each region, and that agents cannot rent their properties. The next iteration of this model should consider these elements in order to become a reliable tool for causal inference and policy advice. Furthermore, policy experiments should also incorporate explicit mechanisms on how the interventions are instrumented as well as precise details on the channels of redistribution. While these limitations are indeed critical to generate reliable policy prescriptions, it is important to remember that none of the existing alternative methods (e.g. regression analysis, equilibrium models, critical analysis, spatial models, etc.) is except of them. Thus, this criticism is general to all models whose use is intended for policy advice.

Overall, the model presented in this paper represents a significant improvement over the existing literature, as it provides a parsimonious way to generate endogenous prices and HWI, while accounting for a high degree of empirical heterogeneity. Furthermore, due to its computational nature, introducing more realistic elements will be easy in comparison to more traditional approaches. Overall, solving the problem of HWI demands holistic paradigms that account for the complexity of the housing market. Hence, computational models such as this one represent the natural direction to be taken towards tackling the staggering inequalities of the 21st century.

Declaration of Competing Interest

The author declares having no conflicts of interest.

Acknowledgements

This work was supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.compenvurbsys.2020.101541

Supplementary data and code to this article can be found at https://github.com/oguerrer/HWI

References

Allegre, G., & Timbeau, X. (2015). *Does housing wealth contribute to wealth inequality? A tale of two New Yorks. Technical report 9, Sciences Po*.

Andrews, D. (2010). Real house prices in OECD countries: The role of demand shocks and structural and policy factors. *OECD Economic Department Working Papers*. *831*. *OECD Economic Department Working Papers* (pp. 5–34).

Andrews, D., & Sánchez, A. (2011). Drivers of homeownership rates in selected OECD countries. *OECD Economics Department Working Papers*. *849*. *OECD Economics Department Working Papers* (pp. 44–). .

Angle, J. (1986). The surplus theory of social stratification and the size distribution of personal wealth. *Social Forces, 65*(2), 293–326.

Appleyard, L., & Rowlingson, K. (2010). Home-ownership and the distribution of personal wealth. *Technical report*. Joseph Rowntree Foundation York.

Arundel, R. (2017). Equity inequity: Housing wealth inequality, inter and intra-generational divergences, and the rise of private landlordism. *Housing, Theory and Society, 34*(2), 176–200.

Atterhög, M. (2005). *Importance of government policies for home ownership rates : An international survey and analysis.* Working Paper 54Stockholm: Royal Institute of Technilogy, Section for Building and Real Estate Economics.

Axtell, R. (2005). The complexity of exchange. *The Economic Journal, 115*(504), F193–F210.

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

Bonnet, O., Bono, P., Chapelle, G., Wasmer, E., et al. (2014). Does housing capital contribute to inequality? A comment on Thomas Piketty's capital in the 21st century. *Sciences Po Economics Discussion Papers*. *7*. *Sciences Po Economics Discussion Papers* (pp. 12–). .

Bowie, D. (2017). *Radical solutions to the housing supply crisis.* Policy Press.

Bulman, M. (18 May 2018). UK facing its biggest housing shortfall on record with backlog of 4m homes, Research shows. *The Independent*.

Chakraborti, A., & Chakrabarti, B. (2000). Statistical mechanics of money: How saving propensity affects its distribution. *The European Physical Journal B - Condensed Matter and Complex Systems, 17*(1), 167–170.

DCLG (2016). *Dwelling stock estimates: 2016, England. Technical report.* Department for Communities and Local Government.

Dewilde, C. (2011). *The interplay between economic inequality trends and housing regime changes in advanced welfare democracies: A new research agenda.* (GINI Discussion Paper).

Dieci, R., & Westerhoff, F. (2012). A simple model of a speculative housing market. *Journal of Evolutionary Economics, 22*(2), 303–329.

Doling, J., & Ronald, R. (2010). Home ownership and asset-based welfare. *Journal of Housing and the Built Environment, 25*(2), 165–173.

Erlingsson, E., Raberto, M., Stefánsson, H., & Sturluson, J. (2013). Integrating the housing market into an agent-based economic model. In A. Teglio, S. Alfarano, E. Camacho-Cuena, & M. Ginés-Vilar (Eds.). *Managing market complexity: The approach of artificial economics, lecture notes in economics and mathematical systems* (pp. 65–76). Berlin, Heidelberg: Springer Berlin Heidelberg.

Erlingsson, E., Teglio, A., Cincotti, S., Stefansson, H., Sturlusson, J., & Raberto, M. (2014). Housing market bubbles and business cycles in an agent-based credit economy. *Economics: The Open-Access, Open-Assessment E-Journal, 8*(2014–8), 1–42.

Feitosa, F., Le, Q., & Vlek, P. (2011). Multi-agent simulator for urban segregation (MASUS): A tool to explore alternatives for promoting inclusive cities. *Computers, Environment and Urban Systems, 35*(2), 104–115.

Filatova, T., Parker, D., & Veen, A. (2009). Agent-based urban land markets: Agent's pricing behavior, land prices and urban land use change. *Journal of Artificial Societies and Social Simulation, 12*, 1–3.

Fisher, L., & Jaffe, A. (2003). Determinants of international home ownership rates. *Housing Finance International, 9*.

Forrest, R. (1995). *Housing & family wealth: Comparative international perspective* (1 ed.). London ; New York: Routledge.

Forrest, R., Murie, A., & Williams, P. (1990). *Home ownership: Differentiation and fragmentation.* London; Boston: Unwin Hyman.

Gallent, N., Durrant, D., & May, N. (2017). Housing supply, investment demand and money creation: A comment on the drivers of London's housing crisis. *Urban Studies, 54*(10), 2204–2216.

Gatti, D., Gallegati, M., & Kirman, A. (2000). Introduction. In D. Gatti, M. Gallegati, & A. Kirman (Eds.). *Interaction and market structure* (pp. 1–7). Berlin, Heidelberg: Springer.

Ge, J. (2014). Who creates housing bubbles? An agent-based study. In S. Alam, & H. Parunak (Eds.). *Multi-Agent-Based Simulation XIV, lecture notes in computer science* (pp. 143–150). Saint Paul, MN, USA: Springer Berlin Heidelberg.

Ge, J. (2017). Endogenous rise and collapse of housing price: An agent-based model of the housing market. *Computers, Environment and Urban Systems, 62*, 182–198.

Geanakoplos, J., Axtell, R., Farmer, J., Howitt, P., Conlee, B., Goldstein, J., ... Yang, C. (2012). Getting at systemic risk via an agent-based model of the housing market. *American Economic Review, 102*(3), 53–58.

Gilbert, N., Hawksworth, C., & Swinney, A. (2009). An agent-based model of the english housing market. *Papers from the 2009 AAAI Spring Symposium, Technical Report SS-09-09* (pp. 30–35). (Stanford, CA, USA).

Harris, J. (22 Jan 2018). We can solve the UK's housing crisis – With a little Imagination. *The Guardian.*.

Henley, A. (1998). Changes in the distribution of housing wealth in Great Britain, 1985-91. *Economica, 65*(259), 363–380.

Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., & Taylor, A. (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics, 134*(3), 1225–1298.

Jordan, R., Birkin, M., & Evans, A. (2012). Agent-based modelling of residential mobility, housing choice and regeneration. *Agent-Based Models of Geographical Systems,* 511–524.

Kouwenberg, R., & Zwinkels, R. (2015). Endogenous price bubbles in a multi-agent system of the housing market. *PLoS One, 10*(6), Article e0129070.

Kraft, H., Munk, C., & Wagner, S. (2018). Housing habits and their implications for life-cycle consumption and investment. *Review of Finance, 22*(5), 1737–1762.

Lersch, P., & Dewilde, C. (2018). Homeownership, saving and financial wealth: A comparative and longitudinal analysis. *Housing Studies, 33*(8), 1175–1206.

Library, H.o. C. (2017). *The housing crisis*.

Lux, T. (2005). Emergent statistical wealth distributions in simple monetary exchange models: A critical review. In A. Chatterjee, S. Yarlagadda, & B. Chakrabarti (Eds.). *Econophysics of wealth distributions: Econophys-Kolkata I, new economic windows* (pp. 51–60). Milano: Springer Milan.

Maxwell, D., & Sodha, S. (2006). *Housing wealth: From first timers to old timers.* London: Institute for Public Policy Research.

McMahon, M., Berea, A., & Osman, H. (2009). *An agent based model of the housing market.* (Working Paper).

OECD (2011). Housing and the economy: Policies for renovation. In OECD (Ed.). *Economic policy reforms 2011 going for growth* (pp. 181–203). .

Pangallo, M., Nadal, J., & Vignes, A. (2019). Residential income segregation: A Behavioral model of the housing market. *Journal of Economic Behavior & Organization, 159*, 15–35.

Parker, D., & Filatova, T. (2008). A conceptual design for a bilateral agent-based land market with heterogeneous economic agents. *Computers, Environment and Urban Systems, 32*(6), 454–463.

Pettifor, A. (27 Jan 2018). Why building more homes will not solve Britain's housing crisis. *The Guardian.*.

Piketty, T. (2014). *Capital in the twenty-first century.* Cambridge MA: Harvard University Press.

Rognlie, M. (2016). Deciphering the fall and rise in the net capital share: Accumulation or scarcity? *Brookings Papers on Economic Activity, 2015*(1), 1–69.

Ronald, R., Kadi, J., & Lennartz, C. (2015). Homeownership-based welfare in transition. *Critical Housing Analysis, 2*(1), 1.

Rowlingson, K. (2002). Private pension planning: The rhetoric of responsibility, the reality of insecurity. *Journal of Social Policy, 31*(4), 623–642.

Ryan-Collins, J. (2019). Breaking the housing–finance cycle: Macroeconomic policy reforms for more affordable homes. Environment and Planning A: Economy and Space, page 0308518X19862811.

Ryan–Collins, J. (2018). *Why Can't you afford a home?* (1 ed.). Cambridge, UK ; Medford, MA: Polity Press.

Ryan-Collins, J., Lloyd, T., & Macfarlane, L. (2017). *Rethinking the economics of land and housing.* London: Zed Books.

Stiglitz, J. (2015). *New theoretical perspectives on the distribution of income and wealth among individuals: Part IV: Land and credit. Working paper 21192.* National Bureau of Economic Research.

Torrens, P. (2001). *New tools for simulating housing choices*.

Weaver, M. (Oct 2017). How did the crisis in UK social housing happen? *The Guardian*.

Wind, B., Lersch, P., & Dewilde, C. (2016). The distribution of housing wealth in 16 European countries: Accounting for institutional differences. *Journal of Housing and the Built Environment,* 1–23.

Yang, F. (2009). Consumption over the life cycle: How different is housing? *Review of Economic Dynamics, 12*(3), 423–443.

Yin, L. (2009). The dynamics of residential segregation in Buffalo: An agent-based simulation. *Urban Studies, 46*(13), 2749–2770.