The emergence of social inequality: a co-evolutionary analysis.

OLIVEIRA, F.S.

2023

The emergence of social inequality: A Co-Evolutionary analysis

Fernando S. Oliveira

Aberdeen Business School, Robert Gordon University, Riverside Building, Garthdee Rd, Aberdeen AB10 7QE, Scotland

A R T I C L E   I N F O

A B S T R A C T

Social inequality is an important issue at the core of economic thought. Based on a stylized pie-sharing game, this article proposes a co-evolutionary computational model to study the interaction between social classes, aiming to answer the following questions: Why do social classes exist, and how do they affect social efficiency and individual effectiveness? Why are wealth inequality and social exclusion persistent? From a methodological perspective, this article extends the pie-sharing game to include a network of interactions and classes, social mobility, evolving wealth, and learning agents. The results show that social inequality is self-emergent. Surprisingly, in the simulations, the existence of social classes increases individual effectiveness, mainly benefiting the poor. Nonetheless, flatter societies (where social classes exist) have a higher average individual effectiveness. As expected, it was observed that wealth inequality is persistent in hierarchical societies, and the upper classes keep a higher proportion of wealth. Furthermore, the analysis extends the bargaining game to include social mobility, showing that, surprisingly, it increases the robustness of the class system and wealth inequality. Finally, the simulation of dual societies shows that these are an evolutionary equilibrium: social exclusion is persistent and accepted by wealthy and poor individuals, and resources are efficiently used.

1. Introduction

Social inequality has been at the core of economic thought since its origins. Both income and wealth inequalities (measured by the Gini Coefficient) have increased since the 1980s in OECD countries. The identified causes of this growing inequality are a decrease in wage shares, the financialization of developed economies with the abandonment of manufacturing, globalization, and a decrease in labor bargaining power (e.g., Piketty, 2014; Arestis, 2018; Tridico and Pariboni, 2018). The importance of reducing inequality is highlighted by studies about developed countries which suggest that inequality leads to lower economic growth and the rise of poverty (e.g., Arestis, 2018, and McKnight, 2019). Nonetheless, this topic is still highly controversial. Dorofeev (2022), using a meta-analysis of previous empirical studies, concludes that an increase (decrease) in inequality is associated with growth when inequality is low (high).

Another stream of research is focused on the existence and definition of social classes. Donni et al. (2015) identify five social classes in the UK based on childhood morbidity, cognitive abilities, parental education, and local deprivation indices. On this same topic, Barozet et al. (2021) present an empirical study on the social stratification in Latin America (Argentina, Chile, and Uruguay) and Europe (Finland, France, Spain, and Great Britain), identifying 3 social classes in Europe, 4 in Chile and 6 in Argentina and Uruguay, respectively. It is clear from the variables used in these studies, such as poverty, education, and health, that they tend to underestimate social divisions, as these are focused on the poorer classes, and there is no measure for extreme wealth inequality and the upper side of the social hierarchy. All these studies have significantly contributed to a better understanding of the social class system. However, as

the social structure is inherently unobservable, it is difficult to estimate social divisions empirically.

For these reasons, the existence of social classes, and their role in society, are important issues that have not yet been fully addressed. This article contributes to the discussion on socioeconomic inequality (captured by initial inequality in wealth distribution) by developing an agent-based model of the pie-sharing game to study how the social structure emerges from the interaction between individuals and affects their behavior, efficient use of resources, and material success (effectiveness). Sharing a pie is a metaphor for situations where different individuals bargain to keep a share of the national income (e.g., Jackson, 2003). The article uses an agent-based model to represent the interactions between different people in society in the context of this pie-sharing game. Agent-based models are a third way of doing science (e.g., Axelrod, 1997), allowing the analysis of non-linear (e.g., Albin and Foley, 1992) and emergent behavior (e.g., Sawyer, 2001, Franke, 2003, Jandhyala and Oliveira, 2021).

From a policy perspective, the critical research questions are related to the rationale for the existence of social classes: Why do they exist, and how do they affect social efficiency and individual effectiveness? What is the explanation for the persistence of wealth inequality and social exclusion?

From a methodological perspective, the contribution is to study how the interactions between initial wealth, social classes, networking, and social mobility, considering reinforcement learning and evolving wealth, affect social efficiency and individual effectiveness.

The article is organized as follows. Section 2 reviews the literature on agent-based models in economic theory. Section 3 describes the extended pie-sharing game and the geometrical forms used to represent the interactions between people in a society. Section 4 describes the agent-based model of the pie-sharing game and the reinforcement learning algorithm. Section 5 simulates how two different societies *solve* the pie-sharing problem. These are the round table (in which everyone has equal standing in society) and hexagonal (in which the number of direct neighbors determines the centrality of a person) societies, with and without social mobility. Finally, Section 6 concludes the article and discusses the robustness of the results and future research avenues.

2. Literature review

The analysis of emergent behavior from learning and adaptation processes has been the object of in-depth study, for example, in game theory (e.g., Macy, 1991; Eguíluz et al., 2005; Centola et al., 2005; Helbing et al., 2005; Ohdaira and Terano, 2011; Bossan et al., 2015). More specifically, Macy (1991) studies the emergence of social norms and cooperation in the Prisoner's Dilemma; Eguíluz et al. (2005) analyze the co-evolution of decisions and social structure in a spatial Prisoner's Dilemma; Centola et al. (2005) model the emergence of social norms using computational learning, studying the conditions under which unpopular norms spread through society; Helbing et al. (2005) explain the emergence of a cooperative equilibrium in a route choice game with congestion; Ohdaira and Terano (2011) show that the complexity of decisions facilitates cooperation in the prisoner's dilemma game, and Bossan et al. (2015) analyze the economic consequences of imitation and social learning in comparison with learning by trial-and-error. For an in-depth discussion on the concept of emergent behavior in sociology and economics, see Sawyer (2001).

This article studies emergent behavior based on the pie-sharing game (e.g., Nash, 1950). As summarized in Ellingsen (1997), this game has many variants, including ultimatum and dictatorship versions (e.g., Poulsen, 2007). Given the multiplicity of existent equilibria, researchers have tried to understand when an equal pie split is the outcome of the game. Axtell et al. (2000) model the emergence of social norms for the distribution of property by assuming that all the agents are randomly paired. They study the impact of introducing differentiation among agents by including exogenous attributes used in the learning process. They analyze the emergence of social norms, i.e., states of the game in which agents' memories and best responses are unchanged, as a function of the exogenous attributes. Oliveira (2010) uses the pie-sharing game to reflect on the interaction between emotions and reason in decision-making under uncertainty. Chiappori et al. (2012) study the two-player pie-sharing game, comparing the stochastic and deterministic versions and providing recommendations on the game's design for empirical studies. Khan (2021) studies the evolutionary stability of the behavioral rules in a pie-sharing game in which each individual interacts with all others, proving that, in equilibrium, each individual claims a time-invariant proportion of the pie. Moreover, the state where individuals claim half of the pie is the only evolutionary stable.

In this pie-sharing game, every outcome in which the pie is completely shared among the players is a Nash equilibrium. The outcomes in which only one person takes the entire pie or the players share less than the total amount are not equilibria. Which factors determine the different equilibria observed? To answer this question, this article uses reinforcement learning to mimic how people learn from their experiences when interacting with others.

There are several different justifications for using reinforcement learning. First, from a psychological perspective, it is well known that reinforcement learning is an essential mechanism in human behavior (e.g., Decker et al., 2016; Kool et al., 2017). Second, in artificial intelligence (e.g., Sutton and Barto, 1998), the reinforcement learning mechanism has been shown to converge to the optimal policy, given a sufficiently large number of learning opportunities. Third, from an economic perspective, reinforcement learning has been used to study several important problems, including market modeling and auction theory (e.g., Bandyopadhyay et al., 2008) and the selection of equilibria in games (e.g., Franke, 2003, Lahkar, 2017, Daskalova and Vriend, 2021).

This article extends the pie-sharing game by including the structure of society and social mobility and by allowing the pie size to depend on how well players have done in previous iterations, i.e., evolving wealth. A basic postulate of this analysis is that each person uses the same strategy to share the pie when interacting with each direct neighbor (e.g., Khan, 2021). Finally, the model uses reinforcement learning (e.g., Sutton and Barto, 1998, Franke, 2003, Daskalova and Vriend, 2021) to represent the behavior of agents in a repeated game.
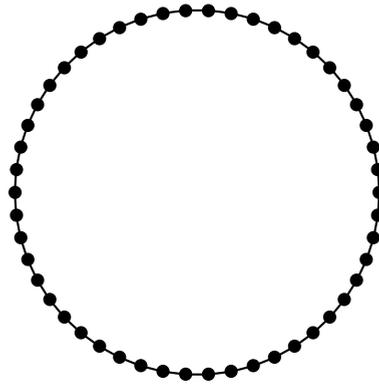
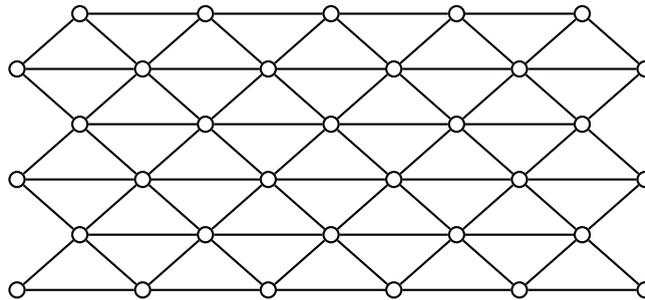**Fig. 1.** Illustration of a round-table society with fifty individuals.

Fig. 2. Illustration of a hexagonal society with thirty-three individuals. It depicts a society with six classes and 5 or 6 individuals per class.

3. The structured bargaining game with social mobility

In the pie-sharing game analyzed in this article, a person i (out of a population of P individuals) maximizes the utility π_i received from the social interactions within a given community of n_i neighbors ($\mathbb{N}_i$). The pie has S slices, and each person i needs to decide how many s_i slices to ask for when interacting with each neighbor. Each individual i has an initial wealth w_i. Both the maximum number of slices S and the decisions ($1 \leq s_i \leq S$) are independent of the size of the pie ($w_i + w_j$).

The utility received by i when playing with the neighbor j, $\mu_i(s_i, s_j)$ is described by equation (1), and can be interpreted as follows: when the proportion demanded by the two players is less than or equal to the size of the pie, each player receives a proportion $\frac{s_i}{S}$ of the total size of the pie, $w_i + w_j$. Otherwise, they receive nothing if they demand more than the available pie size.

$$\mu_i(s_i, s_j) = \begin{cases} \frac{s_i(w_i + w_j)}{S}, & \text{if } \frac{s_i + s_j}{S} \leq 1 \\ 0, & \text{Otherwise} \end{cases} \tag{1}$$

It follows from equation (1) that the utility received by each person depends on the dyadic interactions between the neighbors. There are multiple Nash equilibria in which i and j can agree to share the totality of the pie. In equilibrium, each person i receives the highest average utility from interacting with all the neighbors, π_i^*, by demanding s_i^* slices. The average reward received by i, by choosing action s, is the sum of the utilities obtained in the bargaining game played with each neighbor, divided by the number of neighbors, as summarized in equation (2).

$$\pi_{is} = \sum_{j \in \mathbb{N}_i} \frac{\mu_i(s_i, s_j)}{n_i} \tag{2}$$

The original pie-sharing game is extended by including the social structure, an essential locality feature, i.e., individuals interact only with their neighbors. This structure is unobservable. The model also considers societies organized in social classes. This model of social structure follows Jackson (2003) and does not depend on the individual occupying certain positions. It persists even if the individuals in the different roles change (as analyzed when considering social mobility in Section 5.2).

For these reasons, the social structure is represented by different stylized geometric shapes. Suppose individuals are disposed along a horizontal line. In this case, there is only one social class (if the wealth or revenue is similar across all the individuals), but the position along the line is a differentiating factor. If the line is transformed into a circle, all individuals are in the same social class and have the same societal position. This *round-table* structure, represented in Fig. 1, is the most egalitarian that can be devised.

If there is differentiation among social classes, the most extreme geometry is a vertical line in which each individual is in their social class (i.e., wealth is a direct function of the position along the line). In another possible geometric structure, the triangle, the number of individuals per class depends on the degree of hierarchy (number of layers in the triangle) and the branching factor. This article uses the hexagonal shape to study social hierarchy. There are a few reasons justifying this choice. First, it is more flexible than a triangle or square design, having the advantage of easily changing the degree of hierarchization of society. Second, the hexagonal shape has been used to study the organization of space (e.g., Christaller, 1933), which is very closely related to social structure, as space is one of the resources used for social differentiation, both in the use of productive assets and residential areas (e.g., King, 2020). In this *hexagonal society* (Fig. 2), individuals are arranged in classes, with a maximum of 6 and a minimum of 2 direct connections.

The primary factor differentiating people's influence in society is the number of goods and services they bargain with. This can be the wealth, revenue, salary, or income at stake in any negotiation process. This analysis considers two different regimes to allocate the *initial endowments or wealth*, w_i, to a given person i: a) everyone is equal and has a stake $w_i = 1$, and b) the stake of i is a function of the class to which the individual belongs; in this case, represented by equation (3), the wealth of one individual is inversely defined by the social class to which the person belongs (this parameter is only used within the hexagonal society); L stands for the number of social classes, and l_i is the social class of person $i = 1, \ldots, L$.

$$w_i = L - l_i + 1 \tag{3}$$

From equation (3), in the hexagonal society, wealthier individuals are connected to the rich, whereas the poor interact with the poor. Hence, the locality of social interactions is essential for understanding the impact of structure on individual behavior.

4. Playing with computational reinforcement learning

This section describes how each agent learns to optimize the strategy used in the game based on reinforcement learning. The model uses the *Q*-learning algorithm, e.g., Sutton and Barto (1998), to infer the value created by a given policy, i.e., the Q-values. Based on reinforcement learning, a person i estimates the total reward (q_{is}) received by choosing s_i slices when bargaining with their neighbors, given the profit π_{is} received by choosing this action. For a person i, the maximum expected total wealth is q_i^*.

The algorithm includes two main components: a) to estimate the q_{is}^*, the final wealth associated with a given optimal strategy; and b) To make a decision using the estimated *Q*-values. a) To estimate the value of the proportion claimed by player i in the bargaining process, the q_{is}, based on a simple exponential smoothing of the utilities received in the past.

The estimation of the *Q*-values is performed using equation (4). Note that ":=" represents an updating operator that computes the value of a variable in the next iteration as a function of its current value. The q_{is} is updated using a learning rate α_{is}, of person i, for slice number s_i, taking into consideration the utility forecasting error $\pi_{is} - q_{is}$. When q_{is} underestimates the actual utility, π_{is}, the Q-value is revised upwards, otherwise is revised downwards. Another way to describe equation (4) is as a weighted average between the previous estimate of q_{is} and the observed π_{is}, in which the learning rate α_{is} is the weight put on the observed rewards.

$$q_{is} := q_{is} + \alpha_{is}(\pi_{is} - q_{is}) \tag{4}$$

b) To make a decision using the estimated *Q*-values. This choice of action is probabilistic, with the probability of choosing the best action changing over time. At the start of the simulation, each agent tries different decisions, including those whose expected value is not the maximum. As the simulation proceeds, the probability of choosing the best action increases, and the exploration of alternative actions decreases to zero. For this reason, the *Q*-learning algorithm has two essential parameters: the learning and exploration rates.

The learning rate α_{is} is the proportion of the *Q*-value estimation error used to update the algorithm to predict the *Q*-value for slice s. The algorithm asserts that this learning rate decreases with a person's experience of the state, and decisions, as summarized in equation (5). The δ is the decay of the learning rate, and ρ_{is} is a cumulative variable tracking the number of times a person has tried action s, updated using the function $\rho_{is} := \rho_{is} + 1$. At the start, when the action has not yet been tried, the learning rate is high, and the *Q*-values are adjusted fast. The learning rate is low later in the simulation, and the *Q*-values are only revised marginally.

$$\alpha_{is} = \frac{1}{1 + \delta\rho_{is}} \tag{5}$$

Generally, a rational person chooses the strategy with the highest payoff. Still, as there is a lack of information about the outcomes of the different decisions, reinforcement learning requires experimentation with choices that are not considered the best so that the best long-term policy is inferred over time. For this purpose, at time t, instead of just following what is estimated to be the best policy, a person i has a probability β_t (the exploration rate) of choosing randomly among all possible actions, as calculated in equation (6), in which d is the decay of the exploration rate with the number of iterations (t). This method asserts that when a person is inexperienced, there is much trial and error with all the available options at the start of the process. With the gain of experience, a person tends to choose the best policy, and rarely is any other option chosen.

$$\beta_t = \frac{1}{1 + dt} \tag{6}$$

The algorithm uses time-dependent learning and exploration rates as it allows a broader and less controlled search of the best policy at the start of the simulation and reduces the possibility of significant changes in policy as the simulation progresses. This is true for experimenting with the learning rate and choosing a non-greedy exploration rate policy. The main advantage of the decaying rates is to

Table 1
Agent-based simulation of the pie-sharing game.

For $x = 1$ to X:
   1. Initialization
     a) Initialize network
       i) Round Table.
       ii) Hexagonal.
     b) For all i:

$$w_i = \begin{cases} 1, & \text{if Equal} \\ L - l_i + 1, & \text{if Social class based} \end{cases}$$

$$s_i^* = \arg\max_s q_{is}$$

     c) For all i, initialize the utility values:

$$\pi_i^* = \sum_{j \in \mathbb{N}_i} \frac{\mu_i(s_i^*, s_j^*)}{n_i}$$

   2. Play Game
     For $t = 1$ to T:

     a) Update the exploration rate: $\beta_t = \dfrac{1}{1 + dt}$
     b) Each player chooses the number of slices. For all i:
       $r \sim U[0, 1]$

$$s_i = \begin{cases} \sim U[1, S], & \text{if } r < \beta_t \\ \arg\max_s q_{is}, & \text{otherwise} \end{cases}$$

     c) Compute the equilibrium payoffs. For all i: $\pi_{is} = \sum_{j \in \mathbb{N}_i} \dfrac{\mu_i(s_i, s_j)}{n_i}$
     d) Update the q-values. For all i:

$$\alpha_{is} = \frac{1}{1 + \delta \rho_{is}}, \quad q_{is} := q_{is} + \alpha_{is}(\pi_{is} - q_{is}), \quad \rho_{is} := \rho_{is} + 1$$

     e) If $T > Z$, update the wealth endowment. For all i:
       $w_i = q_i^*$
     f) If there is social mobility:
       Sort the q_i^* in descending order
       Assign each i to the new social position as a function of q_i^*

---

reflect the decrease in the value of experimentation as the agents gain more information, ensuring a reduction of noise with time and experience. The learning and exploration rates are the same for all the agents to ensure that only the social structure influences the results.

When a person chooses the best strategy, the bargaining proportion is s_i^*, as described by equation (7), representing the choice with the highest Q-value.

$$s_i^* = \arg\max_s q_{is} \tag{7}$$

Another significant extension of this article to the pie-sharing game is considering evolving wealth. In this society, the wealth of each person changes over time. The estimated highest wealth, q_i^*, changes over time, as it depends on the strategies used by other agents. For this reason, it needs to be inferred by playing the game. This is because the q_i^* represents the long-term value of the optimal policy, i.e., it is an estimate of the weighted average of all the rewards received by the agent ad infinitum. This estimate gives different weights to the different rewards. These weights exponentially decay with time.

It then follows from equation (4) that both the value at time zero and the value at time T of the wealth is determined by the sequence of profits expected by the individual playing a given strategy.

For this reason, the initial wealth w_i is just an estimate of how the individual will perform in the game. And this represents the endowment received when playing with others.

However, as the game proceeds and the agents learn to optimize their decisions, the expectation about how the agent will perform in the long run is updated and it is equal to the best estimate of i's wealth, i.e., the expected value of all the rewards received by following the optimal policy, as obtained from equation (4), and it is summarized in equation (8).

$$w_i = q_i^* \tag{8}$$

Nonetheless, the wealth remains equal to the initially assigned endowment w_i for a predefined warm-up period (this is because the agents are still learning the best strategy to play.) To motivate the requirement for a warm-up period, consider how people learn in real life. Children and adults are sent to school and educated for decades. Barriers to entry into managerial decisions are implemented so that learning takes place before anyone can manage wealth. The role of the warm-up period is to allow the agents to learn the basic strategy to play the game before any wealth is waged on the decisions. Then, after the warm-up period, at any given iteration t, wealth is equal to the q_i^*.

Table 1 describes the pseudo-code of the game.

The simulation results are analyzed using X different episodes, and Z is the warm-up period. A simulation has P agents who can choose several pie slices to share from 1 to S. Depending on the social structure, the neighborhoods are assigned within a round-table or hexagonal society at the initialization stage. Then, for each agent i, depending on the wealth distribution, w_i is equal to 1 (everyone has
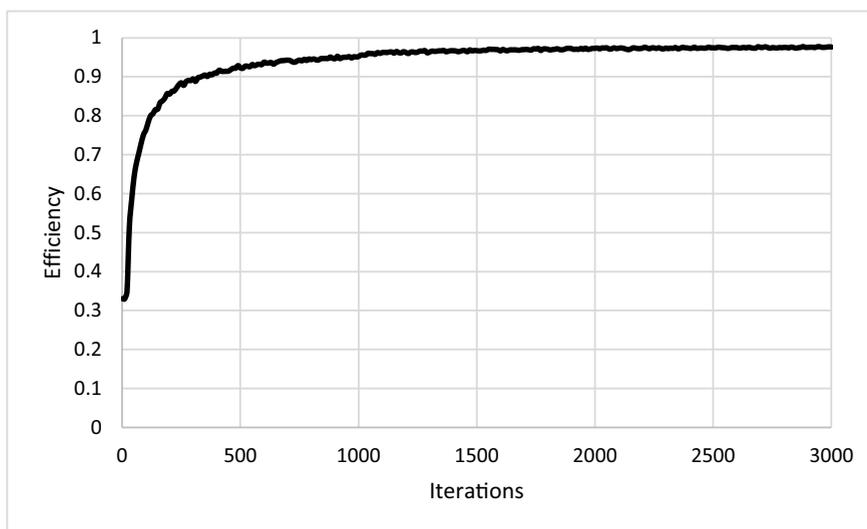
Fig. 3. Social efficiency.

a similar weight in determining the pie size) or a function of $i's$ social class. The initialization stage ends by randomly assigning upper values for the different Q-values and initializing the first choice of the number of pieces of pie demanded by everyone, after which the overall utilities are calculated.

The bargaining game then starts and is repeated for T iterations (long enough for convergence to occur). The game is structured as follows: a) The exploration rate β_t is updated; b) Each agent chooses the number of slices demanded at iteration t. This choice is random with probability β_t (from a discrete uniform distribution) or the action with the highest estimated Q-value is chosen. c) For each i, given the chosen number of pieces of pie demanded, s_i, the total utility received by playing the game is calculated. d) The learning rate for each agent-action pair, α_{is}, is calculated. Then the Q-values and the number of times a given s_i was chosen, ρ_{is}, are updated. e) After the warm-up period, wealth is asserted to be equal to the policy with the highest Q-value, q_i^* for each i. f) When social mobility is considered, the individuals are sorted by descending order of the q_i^* and allocated the respective social positions accordingly.

5. Agent-Based simulation of the structured pie-Sharing game

This section describes the simulations. There are $P = 120$ agents: a multiple of 6, the number of neighbors in the hexagon society, and a significant enough number to derive results that do not change with a higher number of players. The bargaining problem considers that individuals argue about shares of 10%, i.e., $S = 10$ (to keep the coordination loss small). Each episode is run for 3000 iterations (T) to allow for convergence of the learning process. Individual wealth is fixed during the warm-up period, i.e., the first 1000 iterations (Z). The number of simulation runs is $X = 100$, equivalent to running homogeneous simulations with 12,000 individuals each. Each individual's social position remains the same during the 3000 iterations of each simulation. The exception is in the analysis of social mobility, in which individuals are allocated different social positions as the simulation progresses and the respective wealth changes. The same initial parameters are used for each run of the model, so the initial conditions are the same; what changes from run to run are the stochastic processes associated with learning and exploration.

The number of neighbors is $N = 2$ in the round-table society and $N = 2, 3, 4$, and 6 in the hexagonal society. The number of social classes in the *hexagonal society* is $L = 1, 2, 3, 4, 6$, and 12. The *initial wealth* (endowment) of person i is $w_i = 1$ when assuming equality, and $w_i = L - l_i + 1$ when the initial wealth is indexed to their social level (social class). The social structure stratification reflects the narrowing of social interactions to a select group of individuals with identical wealth and the increased stratification of the social system as a function of wealth. Empirical studies tend to capture the latter using the Gini Coefficient (e.g., Arestis, 2018, Tridico and Pariboni, 2018McKnight, 2019); the former is usually not captured empirically.

The learning rate decay (δ) in the base case is set at 0.99, as each decision is visited a limited number of times. The algorithm keeps the learning rate high for as long as possible to slow the convergence rate, avoiding possible local optima. The exploration rate decay (d) is set to 0.05 to slow down convergence to local optima so that the reinforcement learning algorithm tries alternative solutions for a large enough number of iterations.

This analysis is based on the concepts of *efficiency* to represent the proportion of overall resources in society used by the players (unused wealth is wasted), and *effectiveness*, defined as the ratio between the accumulated and initial wealth owned by each individual in society. Moreover, the analysis addresses the impact of networking, wealth (defined as the contribution to the overall size of the *pie* under bargaining), and social mobility on the game's outcome.
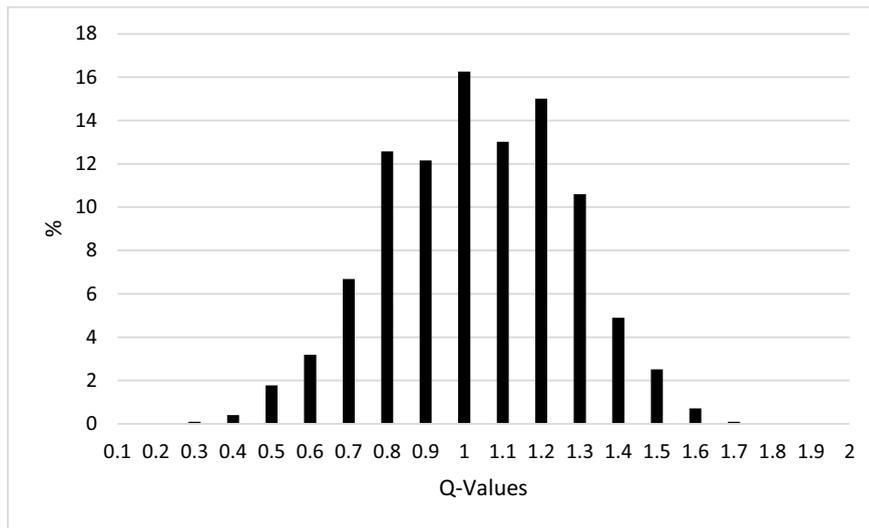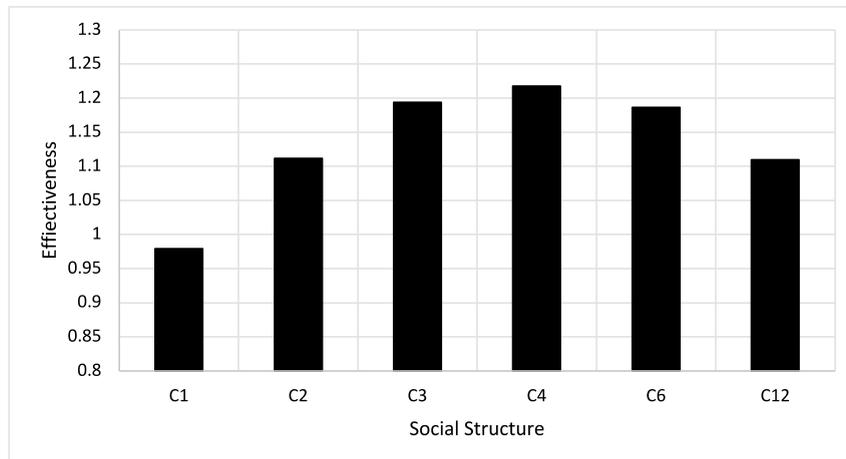
Fig. 4. Q-value distribution.

Fig. 5. Hexagonal society: average individual effectiveness as a function of the social structure.

5.1. Why do social classes exist?

What is the reason for the existence of social classes? How do they interact with or condition individual behavior? Section 5.1.1 addresses this issue by analyzing the round-table society and dummyTXdummy– describing the emergence of social inequality. Then, Section 5.1.2 analyzes how a pre-existent social hierarchy is beneficial or detrimental for the different individuals and social classes.

5.1.1. Round-Table society and the emergence of inequality

This section tests the emergence of social inequality in a society in which all individuals have the same number of neighbors and the same initial wealth. The aim is to test if equality remains or if inequality arises when everyone is equal in every way at the start of the game.

Fig. 3 describes the convergence to equilibrium regarding social efficiency in the round-table society, with about 98% of resources being used (there is only a 2% loss in efficiency).

Fig. 4 summarizes the final wealth distribution for the round-table society with initial equal wealth. Almost all individuals accumulate between 0.4 and 1.5 of their initial endowment as final wealth (Q-value).

Proposition 1. *In a round-table society, wealth inequality is self-emergent.*

From Proposition 1 it follows that a society based on absolute equality is not sustainable in the long term. Eventually, an evolutionary equilibrium (e.g., Maynard-Smith, 1982; Khan, 2021) is reached in which wealth is unequally distributed. This result is in line

Table 2
Explaining Individual Effectiveness for Different Social Structures.

	C2	C3	C4	C6	C12
const	0.016**	0.088***	0.261***	0.555***	0.819***
	(0.041)	(0.005)	(0.003)	(0.003)	(0.002)
N/ω	0.042***	0.015***	0.004	-0.044***	-0.127***
	(0.011)	(0.003)	(0.003)	(0.003)	(0.004)
$1/\omega$	1.296***	1.741***	1.81***	1.763***	1.748***
	(0.000)	(0.014)	(0.011)	(0.011)	(0.014)
$R^2 - Adj.$	0.654	0.872	0.926	0.926	0.842
S. Error	0.265	0.195	0.149	0.129	0.139
N. Observations	12000	12000	12000	12000	12000

Note: In parenthesis are the standard errors. The significance levels are 1% (***) and 5% (**). Non-significant parameters have no asterisk.

Table 3
Mobility Probability between Q-Values (as %).

	1	2	3	4	5	6	7	8	9	10	11	12
1	77	23										
2	10	71	18									
3	3	3	69	25								
4	1	1	5	60	32	1						
5	1	1	2	5	52	36	2					
6	1		3	3	5	45	41	2				
7	1	1	1	3	2	2	43	41	2			
8	1			1	2	1	4	8	39	41	3	
9	1	1	1	2	1	3	2	5	34	48	2	
10	1	1			2	1	1	4	10	38	35	6
11	1		1		1	3	1	2	4	13	41	33
12	1			1	1	3	1	2	4	6	20	60

with Dorofeev (2022), which reports that increased inequality is associated with economic growth in societies with low inequality.

Alternatively, a fully connected graph (in which each individual is connected to everyone else), as in Khan (2021), would also represent a society where all individuals have equal positions. Inequality, in this case, can also emerge as long as the agents have different performances. This is what happens in each one of the neighborhoods of the structured bargaining game analyzed in this article. For this reason, Proposition 1 would also hold in a fully connected graph. This outcome is produced by the different agents' performances emerging while learning to play the game. The additional benefit of having localized interactions is for wealth to be differentiated among individuals and social classes.

5.1.2. The contribution of social classes in a hierarchical society

The analysis now proceeds by focusing on a hierarchical society in which the accumulated wealth is conditioned on the number of social classes, initial wealth, and status (number of neighbors) of each individual. This section investigates what is the role of social classes in human societies by analyzing how they affect individual effectiveness in the pie-sharing game. The study considers six types of hexagonal societies with social structures C1 (one social class) to C12 (12 social classes). The analysis of the average individual effectiveness in Fig. 5 shows that, in the analyzed cases, flatter societies with four classes have the highest individual effectiveness. These results are summarized in Proposition 2. This is in line with the empirical results reported by Barozet et al. (2021), which report that wealthier countries have three or four social classes. Moreover, it also supports the general theory in Dorofeev (2022): the increase of social inequality is associated to higher effectiveness for social structures C1 to C4 and lower effectiveness for societies C6 and C12.

Proposition 2. *Under the specific conditions of the simulation setting, there is an optimum number of social classes (higher than 1) that maximizes individual effectiveness.*

In the hexagonal society, the social structure and the number of neighbors with whom each agent plays the game affect wealth distribution. To better characterize these effects, Table 2 depicts the results of a statistical analysis of the hexagonal society result: it depicts individual effectiveness as a function of the number of neighbors (N) and initial wealth (ω). The C2, C3, C4, C6, and C12 represent the models for the societies with 2, 3, 4, 6, and 12 social classes, respectively. In the hexagonal society with one social class, the statistical evidence shows that the number of neighbors (either one or two) has no significant impact on individual effectiveness: the t-Test is presented in the Appendix, Table A1. The regression models depicted in Table 2 were estimated with 12,000 data points. The $R^2 - Adj.$ ranges from 65.4% to 92.6%, being highest for the societies C4 and C6.

In Table 2, the two explanatory variables are N/ω and $1/\omega$. First, the parameters associated with $1/\omega$ are all significantly positive (at 1% significance level) and above one, showing the higher the initial wealth, the lower the individual effectiveness. Hence, the poorer classes benefit the most from participation in the pie-sharing game.

Second, N/ω captures how the relationship between the number of neighbors and the initial wealth of each individual impacts
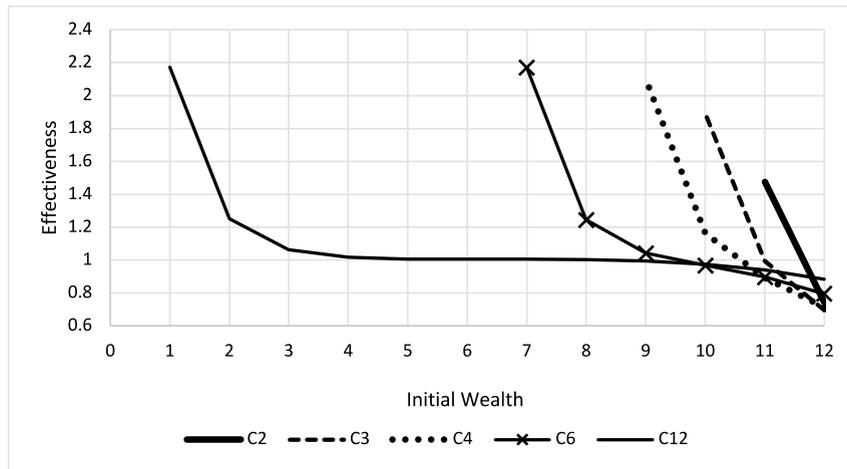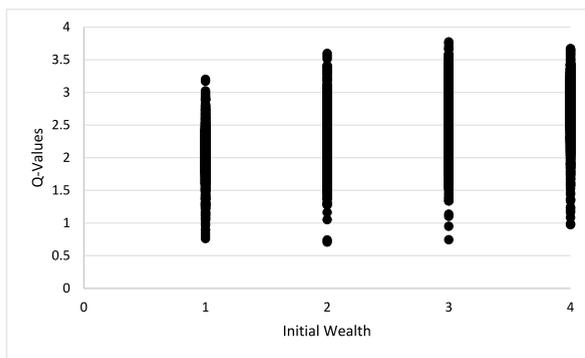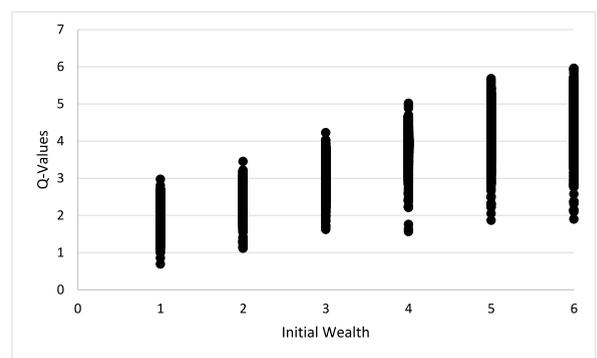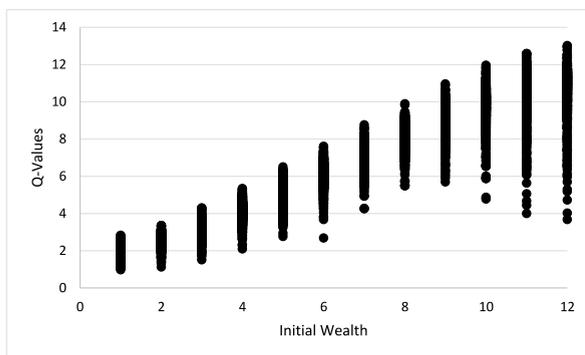
Fig. 6. Average individual effectiveness per social Structure.

(a) Social Structure C4.

(b) Social Structure C6.

(c) Social Structure C12.

Fig. 7. Hexagonal society: impact of wealth and social class on the Q-values.

effectiveness. Most interestingly, the coefficients associated with N/ω range from 0.042 for the C2 society to -0.127 for the C12 society. This result means that in the flatter societies (C2 and C3), the higher the number of neighbors (especially for the poorer individuals), the higher the effectiveness. Consequently, opportunities for social interactions increase individual effectiveness.

On the other hand, for the more hierarchical societies (C6 and C12), the higher number of neighbors, the lower the effectiveness. Hence, social interactions, or opportunities to play the game, represent increased competition and destroy individual effectiveness. The poorer individuals are the most affected by this phenomenon.

Fig. 6 depicts the average individual effectiveness as a function of initial wealth, accounting for the different social structures. Each
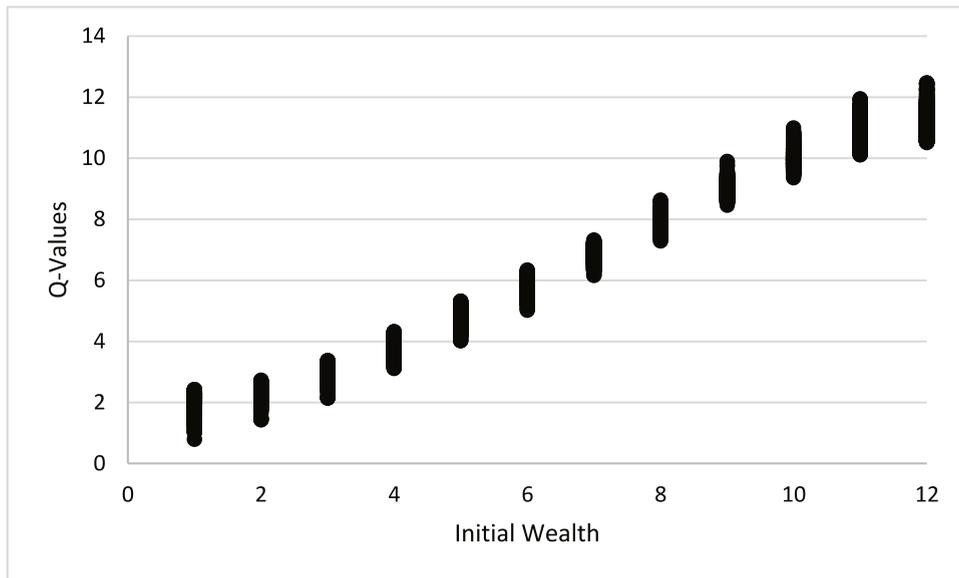
Fig. 8. Hexagonal society with social mobility: impact of initial wealth on the *Q*-values.

line represents the individual effectiveness in the social structures C2 to C12. In flatter societies (C2, C3, and C4), there is a higher transfer of wealth, through the economic mechanism of bargaining, from the upper to the lower classes, as summarized in Proposition 3.

Proposition 3. *The higher the number of social classes and initial wealth differences, the higher the proportion of wealth kept by the wealthier class.*

In hierarchical societies (C6 and C12), there is a more significant transfer of wealth from the upper and middle classes to the lower classes. But overall, the upper class keeps the highest proportion of wealth in the C12 society. Fig. 7 depicts the interaction between social structure, initial wealth, and the Q-values.

In society C4 (showing the highest individual effectiveness), wealth is no longer the determinant of social position in equilibrium. All four social classes have a very similar distribution of wealth (the difference in the mean is insignificant in statistical terms). In this case, the poorer classes managed to recover from their initial disadvantage, resulting in a very equitable society in terms of wealth distribution.

Additionally, it follows from Fig. 7 that, in hierarchical societies C6 and C12, the higher the initial wealth, the higher the average final wealth (*Q*-value). These results are summarized in Proposition 4.

Proposition 4. *(a) In the society with the highest individual effectiveness, equilibrium wealth is independent of social class. (b) In societies with higher initial wealth inequality, the final wealth (Q-value) positively correlates with the initial wealth.*

5.2. Why is social inequality persistent?

So far, the simulations have assumed that in the hexagonal society, the individuals in each position in the social structure remain the same throughout time. This means that each individual keeps playing the game with the same people, independently of performance. However, in reality, when people lose at the game for too long, they also lose their social and (or) business networks. For this reason, in this section, the simulations consider the impact of social mobility on individual wealth.

The agents are sorted from the highest to the lowest q_i^* in each iteration. Then each one is assigned a position in society, from the upper to the lowest class, from left to right: the wealthier agents are placed in higher wealth classes, and the poor are assigned to impoverished social classes.

5.2.1. Impact of social mobility on wealth inequality

Fig. 8 depicts the Q-values of the C12 society with social mobility. The Q-value averages for each class are not statistically different in Figs. and 8, for each class. However, their variances are very different: social mobility decreases the variability of final wealth.

Fig. 8 suggests that when measuring the difference in wealth between classes, the results are biased by the selection factor (who is assigned to each class) and persistent. From a policy perspective, social mobility policies do not destroy the underlying structure of class privilege but instead increase the wealth differences between classes, strengthening the hierarchy. This conclusion is summarized in Proposition 5.

Proposition 5. *Social mobility increases the robustness of the class system and wealth inequality.*
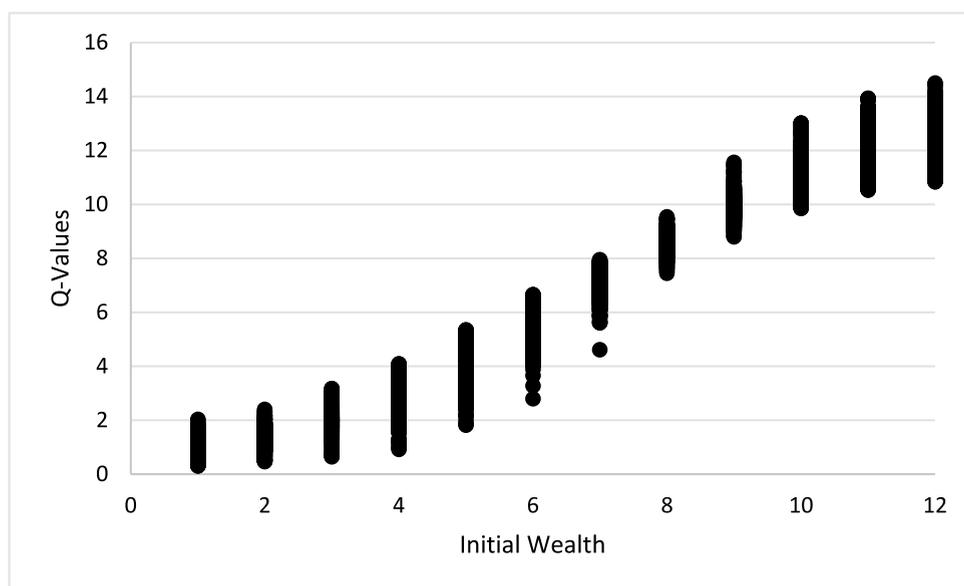
Fig. 9. Dual Society with Social Mobility: Q-values vs. Initial Wealth.

This result raises the question: if social mobility helps to perpetuate wealth and social divisions, why do politicians pursue this policy in the name of social equality? To answer this question, it helps to analyze the impact of social mobility on the probability of moving up or down in the social hierarchy.

The observed transition probabilities between social classes are summarized in Table 3. Each cell in row *r* represents the probability of a person in class *r* reaching the class in column *c* as their final wealth ranking. The probabilities for each row add up to 1. For example, in Table 3, the probability of starting the simulation with a wealth value of 1 and ending with a *Q*-value 1 (2) is 77% (23%). The probability of starting with wealth in the range of 5 to 12 and ending up in the poorest class is 1 percent for each starting wealth value.

On the other hand, each column stands for the breakdown of class members as a function of their original class. For example, in Table 3, the probability of ending up with a Q-value of 1 and starting with a wealth of 1, 2, and 3 is 77%, 10%, and 3%, respectively. Table 3 shows that for each class, the probability of increasing is always higher than decreasing the *Q*-value (except for the wealthiest class). So, from this perspective, social mobility is indeed good for individuals and society.

However, moving up is much harder than moving down: a) only in the class with an initial wealth of 8 are a small percentage (3%) of individuals jumping three classes up; otherwise, for the most, there is a jump of two classes; b) the probability of moving to the poorest classes is the same, for an initial wealth ranging from 4 to 12. These observations are summarized in Proposition 6.

Proposition 6.   *a) The probability of moving up the social hierarchy is higher than that of moving down. b) Social mobility is directionally asymmetric: the size of jumps up the social hierarchy is smaller than the size of falls.*

5.2.2. Asymmetric opportunities in a dual society

Societies where all individuals have the same decision abilities and the same opportunities are, at best, approximated by developed countries. However, societies in developing countries are highly divided, asymmetric, and dualistic, with a privileged group and another group of impoverished individuals (e.g., Boeke, 1953; Singer, 2007).

This section uses the exploration rate to model the access of individuals to knowledge and information in dual societies when choosing their bargaining strategies. The aim of this study is to explore how the asymmetric access to knowledge and information influences individual effectiveness.

This simulation is based on the hexagonal society, with twelve classes, assuming that wealth is a function of social class, as described by equation (3). The exploration rates used in this dual society are fixed and equal to 0.05 (upper three classes), 0.1 (for classes with initial wealth 9 to 7), 0.5 (for classes with initial wealth 6 to 4), and 1.0 (for classes with initial wealth 3 to 1). The final Q-values per initial wealth are depicted in Fig. 9.

By comparing the Q-values in Figs. 8 and 9 it is evident that the wealthier classes benefit the most from the worst access to information and knowledge by the poorer classes. On the other hand, due to their higher exploration rate, the poorest classes, in comparison, are all worse off. Therefore, in these dual societies, there is an incentive for the wealthier classes to keep the poorest classes uneducated and with limited access to information. Nonetheless, the percentage gains of the wealthier classes are small compared to the losses of the poorest classes, as summarized in Proposition 7.

Proposition 7.   *In dual societies, the wealthier classes benefit from the poorest classes' limited access to knowledge, while proportionally the poorest classes lose significantly more.*

Table 4
Mobility Probability between Q-Values (as %).

	1	2	3	4	5	6	7	8	9	10	11	12
1	52	37	12									
2	28	26	31	15								
3	14	18	22	28	17	1						
4	3	10	15	22	36	14						
5	2	5	10	17	19	39	7					
6	1	3	8	13	13	24	35	3	1			
7				2	5	7	41	40	4	1		
8			1	1	4	4	5	41	41	4		
9				1	3	4	3	6	38	40	5	1
10					1	3	3	4	10	38	34	6
11		1		1		1	2	3	4	12	42	32
12		1			1	3	3	3	3	6	18	61

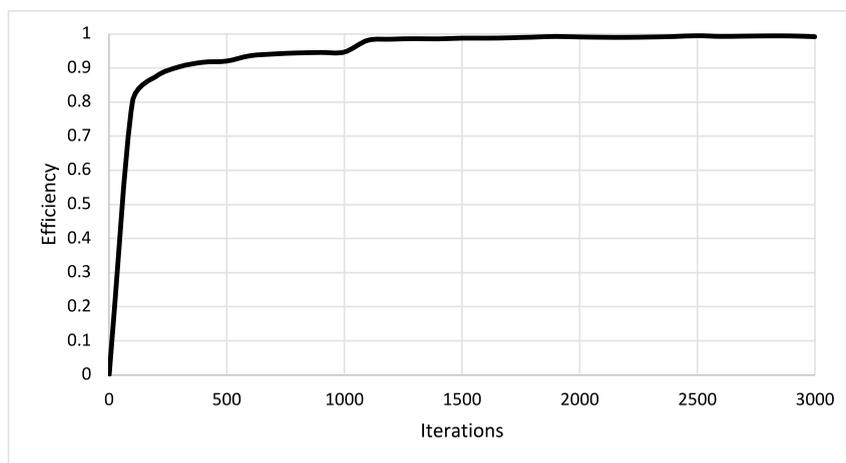**Fig. 10.** Social Efficiency in a Dual Society.

Table 5
Impact of the Dual Society on Individual Effectiveness.

	Mobile	Dual
const	0.84**	0.707***
	(0.002)	(0.012)
N/ω	− 0.178***	0.469***
	(0.003)	(0.020)
$1/\omega$	1.834***	− 1.135***
	(0.011)	(0.091)
$R^2 - Adj.$	0.88	0.095
S. Error	0.11	0.92
N. Observations	12000	12000

Note: In parenthesis are the standard errors. The significance levels are 1% (***) and 5% (**). Non-significant parameters have no asterisk.

Which conditions perpetuate these social structures, and why would the poorest accept this social deal? To answer this question, Table 4 describes the transition probabilities between social classes in a dual society.

Table 4 provides some major insights. There is almost no downward social mobility between the richest and poorest classes, while the probability and size of social jumps are much higher for the poorer classes.

The consequences of these two features of the dual society are as follows: a) for the poorest classes, this society provides the highest possibility of an *effectiveness* increase. (For example, an individual in the poorest class, with starting wealth of 1, has a 12% probability of reaching a Q-value of 3, multiplying their initial wealth by 3, which was found not to be possible in a society with equal exploration and learning rates, i.e., equal opportunities.) b) The upper classes have a much lower risk of moving down to the poorest classes. This is to their significant benefit. Moreover, as seen in Fig. 10, the average social efficiency converged to 1, showing that the dual society achieves an equilibrium in which all the resources are being used, and there is no coordination failure.

These results are summarized in Proposition 8, and are partially in agreement with Axtell et al. (2000): they report that long-lived regimes may emerge which are neither equitable nor efficient. The result in Proposition 8 is stronger: an extremely unequal society is an evolutionary equilibrium in which resource utilization is highly efficient. For this reason, whereas from the perspective of Axtell et al. (2000), inequality is associated with poorer social outcomes and, therefore, undesirable, Proposition 8 proves that social inequality can be the basis of highly efficient societies, providing an argument for the emergence of social classes and the perpetuation of social inequality.

Proposition 8.   *In dual societies, in the evolutionary equilibrium: a) The lower classes have the highest probability of moving up. b) The upper classes have a much lower risk of moving down. c) All resources are efficiently used.*

We proceed with a statistical comparison of the determinants of individual effectiveness in society C12 with social mobility (described in Section 5.2.1) with the dual society. Table 5 summarizes the results.

The C12-mobile and C12-dual societies, despite having the same number of classes, are fundamentally different in how wealth and social interactions explain individual effectiveness. First, social interactions and initial wealth explain individual effectiveness in the C12-mobile society as the $R^2 - Adj. = 88\%$, but not in the C12-dual society as the $R^2 - Adj. = 9.5\%$. Second, even though the model parameters are all significantly different from zero, the respective signs changed. In the C12-mobile society, for any fixed initial wealth, the higher the number of neighbors, the lower the effectiveness due to increased competition. On the other hand, in the C12-dual society, a higher number of neighbors significantly leads to higher performance on average. Third, while from variable $\frac{1}{\omega}$ it follows that the increase in wealth decreases individual effectiveness on average in the C12-mobile society, it *increases* individual effectiveness in the C12-dual society. This means that in the C12-dual society, wealth transfers from the poorer (and less connected individuals) to the wealthier (and better-connected members of society).

6. Conclusion

Francis (2013) and Piketty (2014) have called the attention of politicians and the mass media to the persistent and growing inequality in income distribution. Their work was the primary motivation for this study on the emergence of inequality and the role of social structure in human societies.

From a methodological perspective, the major innovation reported in this article is modeling the pie-sharing game considering social hierarchies, evolving wealth, social mobility, and Q-learning. This article has analyzed an evolutionary bargaining game in which a society is organized under different social structures: a round table (in which social positions are equal) and a hexagon (modeling social hierarchies within a network). This computational model, even though stylized, is complex enough to allow the emergence of surprising and exciting behaviors. The article provides new insights into the existence and persistence of social inequality.

6.1. Why do social classes exist?

This article contributes to policy analysis at different levels. First, it has discussed how inequality is self-emergent. Second, it has shown that individual effectiveness is maximized in flatter societies as the poorer individuals benefit significantly from pie sharing. (Nonetheless, even though the number of social classes is low, individual effectiveness is maximized when there is a degree of social differentiation.) Moreover, increased hierarchic levels protect the upper classes from sharing their wealth with the poor. Most importantly, an increase in the individual effectiveness of the poorer classes results from social inequality: when learning and exploration abilities are independent of social class, the poorest benefit the most from participation in the pie-sharing game.

Additionally, the article has analyzed the impact of social mobility (often seen as a significant policy for improving the condition of the poor) on individual effectiveness. Surprisingly, when learning and opportunities are the same for everyone, social mobility increases the social divide and the perceived differences between classes. In addition, social mobility is directionally asymmetric. The probability of moving up is higher than that of moving down the social hierarchy, but the falls are significantly larger. This means it is relatively easy to increase wealth but very hard to improve significantly. It is easier for someone to lose a substantial proportion of the initial wealth, but a rarer event.

6.2. Why is social exclusion persistent?

The article has also analyzed the outcomes of the bargaining game in dual societies, in which access to information and knowledge is unequal across classes. As observed, wealth flows toward the wealthier, and the poor become even poorer. Most interestingly, however, these dual societies represent an evolutionary equilibrium as the poorer classes are content with their condition (as they have a higher probability of wealth increase within the lower levels of society), and the wealthier have a much higher probability of keeping their wealth. It follows from the analysis of dualistic societies that social exclusion is persistent and accepted by both the wealthier and the poor.

Proposition 8 proved that these dual societies are an evolutionary equilibrium. Is this equilibrium inevitable? Is it the result of a path-dependent process that can be corrected and avoided somehow, or is it independent of the starting conditions? Societies have several mechanisms to regulate redistribution and access to information beyond economic interactions. Differentiated access to information and the high cost of knowledge are central to maintaining these unequal societies. The political process, which is not always

Table A1

t-Test of the impact of the number of neighbors on individual effectiveness.

Number of Neighbors	One	Two
Mean	1.0029	0.9792
Variance	0.0650	0.0569
Observations	200	11800
Hypothesized Mean Difference	0	
df	205	
t Stat	1.3062	
P(T<=t) two-tail	0.1929	
t Critical two-tail	1.9716	

controlled by the dominant social classes, and the willingness of people to change the *status quo* (even those who the system benefits the most) into a fairer society may disrupt the equilibrium. Nonetheless, such a battle, going against the economic forces leading to higher inequality, is difficult to win.

6.3. Discussion

This article uses a stylized model to study the emergence of social inequality, focusing on two geometric forms of social structure: the round table (known to represent the geometry of equals) and a hexagon society (representing a social hierarchy).

The hexagonal society is used to analyze how the local interactions imposed by a social structure affect individual effectiveness and social efficiency. The model was also tested under other geometries, such as triangles and squares. In all the hierarchical societies, independently of the shape, the main qualitative results remained the same: a) the existence of social classes increases individual effectiveness; b) highly hierarchical social structures benefit the upper classes; c) social mobility increases the robustness of the social class system; d) the dual society is evolutionary stable. All these features remain the same because the fundamental property of the social structure, the locality, is similar: only the maximum and minimum numbers of neighbors differ from one social network to another. (The hexagon allows for more complex interactions than the triangle or the square.)

Finally, this article analyzed the emergence of social inequality and the existence of social classes from an economic perspective. Future work may examine the ideas proposed in this article under different games, rules, profit functions, and their evolutionary properties. Additionally, empirical and experimental studies on the existence of social structures and how individuals interact with each other, are also an exciting avenue of research.

Declaration of Competing Interest

None.

Data availability

No data was used for the research described in the article.

Appendix A

Table A1. t-Test of the impact of the number of neighbors on individual effectiveness.

Number of Neighbors	One	Two
Mean	1.0029	0.9792
Variance	0.0650	0.0569
Observations	200	11800
Hypothesized Mean Difference df	0	
t Stat	205	
P(T<=t) two-tail	1.3062	
t Critical two-tail	0.1929	
	1.9716	

References

Albin, P., Foley, D., 1992. Decentralized, dispersed exchange without an auctioneer: a simulation study. J. Econ. Behav. Org. 18, 27–51.

Arestis, P., 2018. Importance of Tackling Income Inequality and Relevant Economic Policies. In: Arestis, P., Sawyer, M. (Eds.), Inequality, international papers in political economy.10.1007/978-3-319-91298-1_1

Axelrod, R., 1997. The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton University Press, Princeton, New Jersey.

Axtell, R., Epstein, J.M., Young, H.P., 2000. The emergence of classes in a multi-agent bargaining model. Center on social and economic dynamics working paper no. 9. february 2000.

Bandyopadhyay, S., Rees, J., Barron, J.M., 2008. Reverse auctions with multiple reinforcement learning agents. Decis. Sci. 39 (1), 33–63.

Barozet, E., Boado, M., Marqués-Perales, I., 2021. the measurement of social stratification: comparative perspectives between Europe and Latin America. In: López-Roldán, P., Fachelli, S. (Eds.), Towards a comparative analysis of social inequalities between europe and latin america.10.1007/978-3-030-48442-2_6 Boeke, J.H., 1953. Economics and Economic Policy of Dual Societies. Tjeenk Willink & Zoon, Haarlem.

Bossan, B., Jann, O., Hammerstein, P., 2015. The evolution of social learning and its economic consequences. J. Econ. Behav. Org. 112, 266–288.

Centola, D., Willer, R., Macy, M., 2005. The emperor's dilemma: a computational model of self-enforcing norms. Am. J. Sociol. 110 (4), 1009–1040.

Chiappori, P.-A., Donni, O., Komunjer, I., 2012. Learning from a piece of pie. Rev. Econ. Stud. 79 (1), 162–195.

Christaller, W., 1933. Die zentralen Orte in süddeutschland. Gustav Fischer, Jena (English Translation: Central Places in Southern Germany. Prentice Hall, Englewood Cliffs. 1966

Daskalova, V., Vriend, N.J., 2021. Learning frames. J. Econ. Behav. Org. 191, 78–96.

Decker, J.H., Otto, A.R., Daw, N.D., Hartley, C.A., 2016. From creatures of habit to goal-directed learners: tracking the developmental emergence of model-based reinforcement learning. Psychol. Sci. 27 (6), 848–858.

Donni, P.L., Rodríguez, J.G., Dias, P.R., 2015. Empirical definition of social types in the analysis of inequality of opportunity: a latent classes approach. Soc. Choice Welf. 44, 673–701. https://doi.org/10.1007/s00355-014-0851-6.

Dorofeev, M.L., 2022. Interrelations between income inequality and sustainable economic growth: contradictions of empirical research and new results. Economies 10, 44. https://doi.org/10.3390/economies10020044.

Eguíluz, V.M., Zimmermann, M.G., Cela-Conde, C.J., 2005. Cooperation and the emergence of role differentiation in the dynamics of social networks. Am. J. Sociol. 110 (4), 977–1008.

Ellingsen, T., 1997. The evolution of bargaining behavior. Q. J. Econ. 112 (2), 581–602.

Francis, P., 2013. Apostolic exhortation evangelii gaudium. Vatican Press.

Franke, R., 2003. Reinforcement learning in the el farol model. J. Econ. Behav. Org. 51, 367–388.

Helbing, D., Schonhof, M., Stark, H.-U., Holyst, J.A., 2005. How individuals learn to take turns: emergence of alternating cooperation in a congestion game and the prisoner's dilemma. Adv. Complex Syst. 8 (1), 87–116.

Jackson, W.A., 2003. Social structure in economic theory. J. Econ. Iss. 727–746.

Jandhyala, S., Oliveira, F.S., 2021. The role of international anti-corruption regulations in promoting socially responsible practices. J. Econ. Behav. Org. 190, 15–32.

Khan, A., 2021. Evolutionary stability of behavioural rules in bargaining. J. Econ. Behav. Org. 187, 399–414.

King, L.J., 2020. Central Place Theory. Web book of regional science. Regional Research Institute, West Virginia University.Edited by Grant Ian Thrall, 1985; Randall Jackson

Kool, W., Gershman, S.J., Cushman, F.A., 2017. Cost-benefit arbitration between multiple reinforcement-learning systems. Psychol. Sci. 28 (9), 1321–1333.

Lahkar, R., 2017. Equilibrium selection in the stag hunt game under generalized reinforcement learning. J. Econ. Behav. Org. 138, 63–68.

Macy, M.W., 1991. Learning to cooperate: stochastic and tacit collusion in social exchange. Am. J. Sociol. 97 (3), 808–843.

Maynard-Smith, J., 1982. Evolution and the theory of games. Cambridge University Press, Cambridge.

McKnight, A., 2019. Understanding the Relationship between Poverty, Inequality and Growth: A Review of Existing Evidence. London school of economics and political science. CASEpaper 216/LIPpaper 8. ISSN 1460–5023.

Nash, J., 1950. The bargaining problem. Econometrica XVIII 155–162.

Ohdaira, T., Terano, T., 2011. The diversity in the decision facilities cooperation in the sequential prisoner's dilemma game. Adv. Complex Syst. 14 (3), 377–401.

Oliveira, F.S., 2010. Modeling emotions and reason in agent-based systems. Comput. Econ. 35, 155–164.

Piketty, T., 2014. Capital in the twenty-First century. The Belknap Press of Harvard University Press, Cambridge Massachusetts, London England.

Poulsen, A.U., 2007. Information and endogenous first mover advantages in the ultimatum game: an evolutionary approach. J. Econ. Behav. Org. 64, 129–143.

Sawyer, R.K., 2001. Emergence in sociology: contemporary philosophy of mind and some implications for sociological theory. Am. J. Sociol. 107 (3), 551–585.

Singer, H., 2007. Dualism revisited: a new approach to the problems of the dual society in developing countries. J. Dev. Stud. 7 (1), 60–75.

Sutton, R.S., Barto, A.G., 1998. Reinforcement-learning: An introduction. MIT Press.

Tridico, P., Pariboni, R., 2018. Theoretical and Empirical Analyses of the Rise of Income Inequality in Rich Countries. In: Arestis, P., Sawyer, M. (Eds.), Inequality, international papers in political economy.10.1007/978-3-319-91298-1_1