



UvA-DARE (Digital Academic Repository)

How social reinforcement learning can lead to metastable polarisation and the voter model

Meylahn, B.V.; Meylahn, J.M.

DOI

[10.1371/journal.pone.0313951](https://doi.org/10.1371/journal.pone.0313951)

Publication date

2024

Document Version

Final published version

Published in

PLoS ONE

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Meylahn, B. V., & Meylahn, J. M. (2024). How social reinforcement learning can lead to metastable polarisation and the voter model. *PLoS ONE*, *19*(12), Article e0313951. <https://doi.org/10.1371/journal.pone.0313951>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

RESEARCH ARTICLE

How social reinforcement learning can lead to metastable polarisation and the voter model

Benedikt V. Meylahn^{1*}, Janusz M. Meylahn²¹ Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Amsterdam, The Netherlands,² Department of Applied Mathematics, University of Twente, Enschede, The Netherlands* b.v.meylahn@uva.nl

Abstract

Previous explanations for the persistence of polarization of opinions have typically included modelling assumptions that predispose the possibility of polarization (i.e., assumptions allowing a pair of agents to drift apart in their opinion such as repulsive interactions or bounded confidence). An exception is a recent simulation study showing that polarization is persistent when agents form their opinions using social reinforcement learning. Our goal is to highlight the usefulness of reinforcement learning in the context of modeling opinion dynamics, but that caution is required when selecting the tools used to study such a model. We show that the polarization observed in the model of the simulation study cannot persist indefinitely, and exhibits consensus asymptotically with probability one. By constructing a link between the reinforcement learning model and the voter model, we argue that the observed polarization is metastable. Finally, we show that a slight modification in the learning process of the agents changes the model from being non-ergodic to being ergodic. Our results show that reinforcement learning may be a powerful method for modelling polarization in opinion dynamics, but that the tools (objects to study such as the stationary distribution, or time to absorption for example) appropriate for analysing such models crucially depend on their properties (such as ergodicity, or transience). These properties are determined by the details of the learning process and may be difficult to identify based solely on simulations.

OPEN ACCESS

Citation: Meylahn BV, Meylahn JM (2024) How social reinforcement learning can lead to metastable polarisation and the voter model. PLoS ONE 19(12): e0313951. <https://doi.org/10.1371/journal.pone.0313951>

Editor: William Ott, University of Houston, UNITED STATES OF AMERICA

Received: June 14, 2024

Accepted: November 2, 2024

Published: December 17, 2024

Copyright: © 2024 Meylahn, Meylahn. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The datasets used to generate Figs 2 and 3, along with the core code for the simulation, are available on GitHub at <https://github.com/Benepher/Data-for-social-reinforcement-learning>.

Funding: This research was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 945045, and by the NWO Gravitation project NETWORKS under grant no. 024.002.003. The funders had no role in study design, data collection and analysis,

1 Introduction

Since at least 1964 scientists have been trying to answer the question “what on earth must one assume to generate the bimodal outcome of community cleavage studies” [1, p. 153]. Possible answers to this question have been presented; bounded confidence [2–5] whereby agents stop listening to others if their opinion is too different from their own, repulsive forces between agents [6–10] based on possible negative connections within a network or messages eliciting the opposite effect within a recipient, stubbornness of an agent toward changing their opinion [11–13], and distinguishing between an agent's expressed opinion and their internal opinion [14, 15]. What unifies these explanations is that the resulting models all include some element from which one might infer the possibility of polarization. Models with only attractive forces

decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

on the other hand, typically lead only to consensus (see for example the models of weighted averaging [16–19] and imitation [20, 21]). Thus it comes as a surprise that the model studied by Banisch and Olbrich [22] with only attractive forces between agents, seems to exhibit persistent polarization.

The dynamics in models of opinion formation typically take place on a network. A network consists of nodes (representing agents) and edges between them (representing social influence or ties). Nodes that share an edge are said to be neighbours in the network. A well studied class of opinion dynamics models on networks from the sociophysics literature is the class of voter models [20, 21] (see [23] for an introduction). In these models a random agent is selected each round to update their opinion. The agent does this by copying the opinion held by one of their neighbours. Reinforcement learning is a model for learning by feedback: actions (or opinions in our case) for which an agent receives positive feedback are reinforced. Actions that receive negative feedback on the other hand are less likely to be taken in the future.

For convenience, in the remainder we refer to the paper of Banisch and Olbrich [22] as BO while we refer to the reinforcement learning model they study as the ‘Asymmetric Reinforcement Learning for Opinion Dynamics model,’ or simply the ARLOD model. This influential model includes no repulsive element in the interaction between agents. It proposes modelling the evolution of opinions (of agents on connected networks) using multiagent reinforcement learning, where agents interact via a coordination game. They find, using simulations, that allowing agents to learn their opinion through trial and error gives rise to the emergence of persistent polarization. This is surprising, because in this model after an interaction between two agents, the opinions of the two necessarily get closer together and cannot remain unchanged or get further apart. That is, in the ARLOD model there are no repulsive forces or assumptions of bounded confidence. Models of opinion dynamics may be classified into ‘assimilative,’ ‘repulsive’ and models with ‘similarity bias’ [24]. The model under consideration here does not traditionally fall in the category of models with only assimilative forces between agents because it utilizes experience based learning. Note that in the original article BO [22], ‘persistent’ and ‘stable’ are used interchangeably. In order to avoid confusion, we use ‘persistent’ to discuss their claims about the ARLOD model and ‘stable’ when making our own claims. Törnberg *et al.* [25] build on the ideas of the ARLOD model by incorporating the role of agent identity. Törnberg [26], similarly looking for drivers of polarisation without the assumption of negative influences but dissatisfied with BO’s assumption of selective exposure (a fixed and constant network), analysed a model which includes non-local interaction to model the effect of media. A variant of the reinforcement learning model with multiple opinions and synchronous updating has been studied in [27]. Their results highlight the difficulty of reaching consensus in complex networks using reinforcement learning. The idea of modelling opinion dynamics by reinforcement learning has been built on since (e.g. [28–31]).

An overarching goal in this paper is to highlight the importance of the relationship between model assumptions and characteristics. It can be tempting to design a model and study its characteristics by simulation. However, to accurately present the results of such a simulation study it may be important to first identify certain model characteristics. We illustrate the importance of this by presenting three results on the ARLOD model by BO [22].

We show analytically that consensus is reached in the ARLOD model with probability one in the long run. The polarisation found in [22] necessarily gives way to consensus eventually. To elucidate this result, we run simulations to estimate the tail probabilities for the time to consensus. We find that these exhibit heavy tails, indicating that there may be metastable states (corresponding to polarization) in which the model resides for a long time before reaching consensus.

The phenomenon of metastable polarization together with eventual consensus has previously been observed in the context of the voter model, first introduced in [21]. Specifically, consensus is reached eventually if the state space of the model is finite [32], and it has been shown that polarisation is metastable in certain network topologies [33–36]. Recently, metastable opinion polarisation has been identified in [37] where it is shown to arise from biased information processing.

The dynamics of the voter model on networks consists of agents adopting one of their neighbours' opinions at random. At first sight, dynamics of this kind seem rudimentary in comparison to the sophisticated dynamics of reinforcement learning. However, we show that, under a separation of time scales, the ARLOD model converges in distribution to a voter model. This relationship highlights that the polarisation observed in the ARLOD model may indeed be *metastable* depending on the network structure. It also bridges the seemingly disparate approaches to modelling opinions: sociophysics and computational sociology. These two approaches differ in their typical level of abstraction, and whether they aim for tractability by keeping the dynamics simple or aim to approach realism by modelling the agents with a relatively high level of sophistication. The ARLOD model falls in the class of computational sociology seeing as the agents in the model are sophisticated enough to learn from experience. The relationship we show between this model and the voter model (a very simple model where agents imitate one another) is thus an example of a bridge between the two approaches to studying polarization.

In designing their model, BO [22] decide to make the interaction-learning relationship asymmetric: only one of the agents partaking in the interaction is allowed to explore and learn from the experience. We show that adapting the model to be symmetric fundamentally changes the nature of the opinion dynamics from being non-ergodic to being ergodic. Under this model, consensus is no longer absorbing so that the tools appropriate for studying polarization and consensus differ from those required in the case of the ARLOD model. For example, in an ergodic system the stationary distribution may be estimated by studying the mean return time to polarized (or consensus) states. On the other hand for a non-ergodic system with absorbing states one typically studies the time to absorption (in a consensus or polarized state if these are indeed absorbing) or the number of visits to transient states before absorption. If there are both consensus and polarized absorbing states, the relative probability of consensus or polarization can be studied.

2 Results

In this section, we present the asymptotic analysis of the asymmetric reinforcement learning for opinion dynamics (ARLOD) model presented by Banisch and Olbrich [22] in the long-time limit, its relation to the voter model and the asymptotic analysis of a symmetric modification of the model.

All three analyses (on the ARLOD model, the symmetrized version thereof, and the relationship between the ARLOD model and the voter model) consider the same reinforcement learning method, namely, Q-learning. By using Q-learning, agents assign an estimate of the “quality” of expressing each opinion to a randomly selected neighbour called a Q-value. We present the ARLOD model for completeness of the current text. We refer to this model as the *asymmetric* model because in the interaction between two agents the roles are distinguishable. One agent is chosen to express their opinion to another, who merely responds. Only the first agent updates their Q-values, and only the first agent can explore.

Different notions of stability exist in various fields related to the model we study. To avoid confusion we present the definition of a stable state as used in this and the Methods section of

the paper. The notion we use is strongly related to the notion of absorption. In the Introduction and Discussion sections we revert to using ‘absorbing’ and ‘stable’ separately as we discuss these notions outside of the paradigm of the model we study here.

Definition 1 (Stability). *A state (or class of states) is stable if once the process has entered this state (class of states), it remains there indefinitely.*

2.1 Asymptotic behaviour

2.1.1 The asymmetric reinforcement learning opinion dynamics (ARLOD) model.

This model of learning through social feedback considers $N \in \mathbb{N}$ agents on a random (connected) geometric network topology [38]. In particular, the network is given by $G = (V, E)$ where V are the vertices representing agents, and E are the connections between agents. The graph is constructed according to the random geometric graph model with radius r_g (for details, see §4.2 and Appendix E). Initially, all agents $i \in \{1, \dots, N\}$ assign a (possibly random) quality $Q_o^i \in [-1, 1]$ to each opinion $o \in \{-1, 1\}$. Note that in the simulation we initialize these values in $[-0.5, 0.5]$ instead of $[-1, 1]$ which is all that is required for the theoretical analysis. We do this following BO’s original simulation. The reason provided is to have on average half the agents favouring each opinion. An agent holds the opinion which they assign the higher quality. In each discrete time step t , an agent i is chosen uniformly at random to express their opinion $o_i(t)$ to a randomly selected neighbour j . This neighbour responds by either punishing them if the expressed opinion differs from their own ($R_j = -1$), or rewards them if the expressed opinion is shared ($R_j = 1$).

Agents thus learn the value of each of the two possible opinions $\{-1, 1\}$ from their experiences using stateless Q-learning. This means that each opinion o is assigned a Q-value Q_o , measuring its “quality”, which is updated as follows for the opinion $o_i(t)$ expressed in round t :

$$Q_{o_i(t)}^i(t + 1) = Q_{o_i(t)}^i(t)(1 - \alpha) + R_j\alpha. \tag{1}$$

Here $\alpha \in (0, 1)$ is called the learning rate. The Q-value of the opinion they did not express is not altered so that

$$Q_{-o_i(t)}^i(t + 1) = Q_{-o_i(t)}^i(t). \tag{2}$$

We assume that the agent chosen to express their opinion exploits their favoured opinion (the one with the greater Q-value) with probability $1 - \epsilon$ and explores by expressing their disfavoured opinion with probability $\epsilon > 0.1$. This is known as ϵ -greedy Q-learning with fixed exploration rate ϵ .

The dynamics per round are depicted in a schematic in Fig 1. Note that only agent i adjusts their Q-values after such an interaction, and that agent j ’s response is deterministic (honest).

2.1.2 Asymptotic consensus and non-ergodicity. We now prove that in the ARLOD model the long-time limit of the dynamics necessarily results in consensus and does not allow for polarization. The proof is inspired by the proof of an analogous result for agents who learn by simple exponential smoothing in [39]. We explore the time to consensus by means of simulation. For the details on the simulation, see §4.1.

Analytical results. Our first result states that consensus is a stable state. In this regard, we define consensus as the state of the model in which the Q-values each agent assigns to the opinions have the same ordering. Note that we use a slightly different notation to that used by BO. We define $Q_o^i(t)$ as the Q-value that agent $i \in \{1, 2, \dots, N\}$ assigns to opinion $o \in \{-1, 1\}$ at time $t \in \mathbb{N}$.

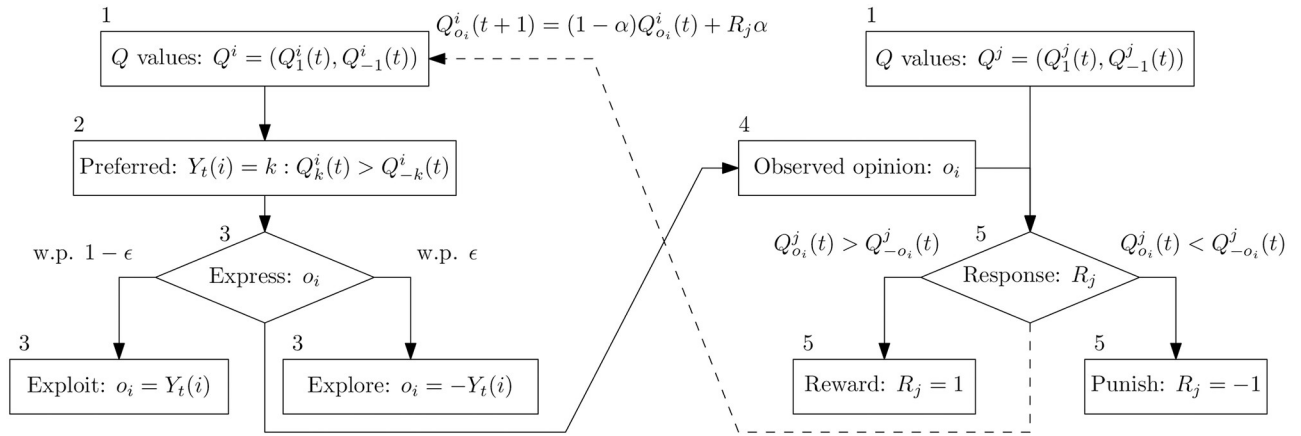


Fig 1. A schematic of the procedure followed by the two agents selected to interact in one round of the ARLOD model, as originally described in [22]. Agent i expresses an opinion O_i to their neighbour j , who responds by punishing or rewarding agent i . Agent i updates the Q-value for the opinion they expressed accordingly. The numbers to the top left of the boxes indicate the suggested order for reading the schematic.

<https://doi.org/10.1371/journal.pone.0313951.g001>

Lemma 1 (Consensus is stable). *If there exists a time t_0 such that $Q_o^i(t_0) > Q_{-o}^i(t_0)$ for some opinion $o \in \{-1, 1\}$, and each agent $i \in \{1, 2, \dots, N\}$, then $Q_o^i(t) > Q_{-o}^i(t)$ for all $t \geq t_0$ and for all agents $i \in \{1, 2, \dots, N\}$.*

We prove Lemma 1 in Appendix B. The proof follows from the fact that agents respond honestly, so that once all agents have the same ordering of Q-values, each exploration is punished while each exploitation is rewarded. This preserves the Q-value ordering.

The next result required to prove that consensus is reached with probability one in the long-time limit, is that consensus is reachable from any state that is not consensus.

Lemma 2 (Consensus is reachable from all other states). *If the learning rate $\alpha > 0$, the exploration rate $\epsilon > 0$, and G is connected then the probability of reaching consensus in finite time is positive, i.e.,*

$$\mathbb{P}(\exists t_1 < \infty : Q_o^i(t_1) > Q_{-o}^i(t_1), \forall i \in \{1, \dots, N\}) > 0, \tag{3}$$

for some $o \in \{-1, 1\}$.

Lemma 2 is proved in Appendix B and hinges on the realisation that the ordering of an agent’s Q-values may switch in a finite number of rounds as long as they have a neighbour whose Q-value ordering differs from theirs. The number of rounds required for this switch to occur is bounded from above by $2r + 2$ with

$$r = \left\lceil \frac{\log(\xi)}{\log(1 - \alpha)} \right\rceil, \tag{4}$$

for some $\xi \in (0, \alpha)$. Note that Lemma 2 is true for all connected graphs between $N < \infty$ agents and all starting states (Q-values of agents) that are not in consensus. Furthermore, consensus on either of the two opinions is reachable in this way.

We now state the first main theorem of the paper, which states that consensus is reached with probability one in the long run in the ARLOD model.

Theorem 1 (Consensus is guaranteed). *If the learning rate $\alpha \in (0, 1)$, the exploration rate $\epsilon \in (0, 1)$, and G is connected then the probability of consensus in the long run is one, i.e.,*

$$\mathbb{P}(\exists t_0 < \infty : Q_o^i(t) > Q_{-o}^i(t), o \in \{-1, 1\}, \forall i \in \{1, \dots, N\}, \forall t \geq t_0) = 1. \tag{5}$$

Proof. By Lemma 1 consensus is stable and so once it is reached it persists. By Lemma 2 the probability of reaching consensus from not consensus in $R = (N - 1)(2r + 2)$ rounds is bounded from below by $p > 0$. Thus, the probability of not reaching consensus in kR rounds is bounded from above by

$$\mathbb{P}(\nexists t_0 \leq kR : Q_o^i(t) > Q_{-o}^i(t), o \in \{-1, 1\}, \forall i \in \{1, \dots, N\}, \forall t \geq t_0) \leq (1 - p)^k. \tag{6}$$

The probability of never being absorbed is then bounded from above by the limit of (6) as $k \rightarrow \infty$ which is zero. Therefore, the probability of the complement is one.

This implies that the polarisation observed as persistent in the presentation of the original model’s simulation cannot persist indefinitely. In particular, the probability reported in Fig 5 of BO should be reinterpreted from ‘probability of consensus’ to ‘probability of consensus before time $N \times 20000$.’ Furthermore, this implies that the probability of the system being in a polarised state tends to zero as $t \rightarrow \infty$.

Note that the conditions on the network are only that it is connected. This is not a significant limitation. Studying polarization is most interesting in connected networks where there is still interaction between agents that disagree. The results also hold separately for each component of a disconnected network. Though consensus within each component does not imply consensus between components.

Simulations. In light of Theorem 1, we investigate the time to consensus as a function of the radius of the geometric network structure by simulation. The parameter settings are stated and motivated in §4.1. We define the time to consensus τ as

$$\tau := \min\{t : Q_o^i(t) > Q_{-o}^i(t), o \in \{-1, 1\}, \forall i \in \{1, \dots, N\}\}. \tag{7}$$

In Fig 2(A), we show the tail probabilities of the time to consensus $\mathbb{P}(\tau \geq t)$ for different radii of the random geometric graph model on a logarithmic scale. A clear pattern emerges; the bigger the radius, the sooner consensus is reached. We also note that the distributions exhibit heavy tails, especially for the smallest three settings of the radius: $r_g \in \{0.25, 0.3, 0.35\}$. This can be seen by the near linear lines (on the log-log scale) which are representative of power-law and log-normal distributions.

In Fig 2(B), we show box and whisker diagrams of the simulated time to consensus (conditioned on $\tau < t_{\max} = 10^{10}$). This representation of the simulated data clearly shows that there are many runs which might be identified as ‘outliers.’ This indicates that the time to consensus has a high skewness and, like the tail probabilities, points towards a heavy-tailed distribution. A possible explanation for the heavy-tails is the existence of metastable states, which the system may spend a lot of time in before eventually ‘jumping’ out to consensus. Indeed, similar heavy-tailed survival probabilities were observed for the voter model on small-world networks, which exhibit metastable polarisation [35]. We see that as the radius r_g decreases, the probability that consensus is reached after time $t \in \mathbb{R}$ increases. This shows how quantitatively the dynamics do depend on the realisation of the network structure.

To illustrate this phenomenon of metastability, we plot the state of the system at different points in time for a single trajectory. In Fig 3 we show the total number of agents holding opinion $o = 1$ over time in this trajectory, which illustrates the metastable behaviour. In Fig 4 we show the network of agents coloured according to their opinion at different time steps. Note that because we select a run which illustrates metastable polarization, the network depicted

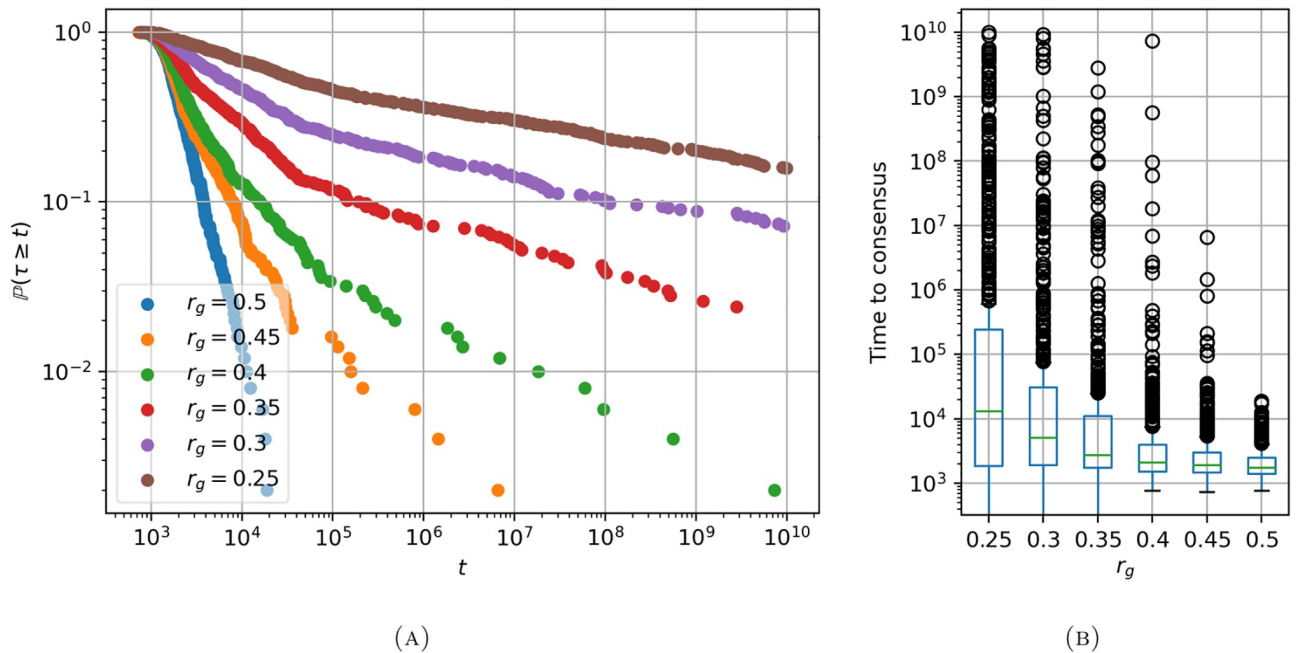


Fig 2. (A) Tail probabilities ($\mathbb{P}(\tau > t)$) (on a log-log scale) and (B) a box and whisker diagram for the time to consensus for different values of the radius r_g used in the random geometric graph model to sample networks. The linear nature of these plots are indicative of a heavy tailed distribution. The high number of outliers on the upper end of the time to consensus is indicative of a heavy-tailed distribution. The parameter settings are detailed in §4.1.

<https://doi.org/10.1371/journal.pone.0313951.g002>




Fig 3. Number of agents holding opinion $o = 1$ in a simulation run exhibiting metastable behaviour plotted with time on a logarithmic scale. The state of the network is plotted for telling timestamps of this simulation run in Fig 4. In this simulation run $r_g = 0.25$, the other parameters are as in §4.1.

<https://doi.org/10.1371/journal.pone.0313951.g003>

here has more community structure than what may be typical of the random geometric network algorithm. This is because the presence of community structure allows nodes to have more in-community than out-community connections and so hold the opinion of their community for a long time (metastable), even if this is not uniform across communities. We see that by $t = 10^4$ two groups emerge; just less than 20 agents holding opinion $o = -1$ and the rest holding opinion $o = 1$. This remains the case until shortly after time step 5.82×10^6 when the opinions all quickly converge to $o = 1$. The long time spent around one state with many small fluctuations followed by a quick exit to a stable state is typical of metastability.

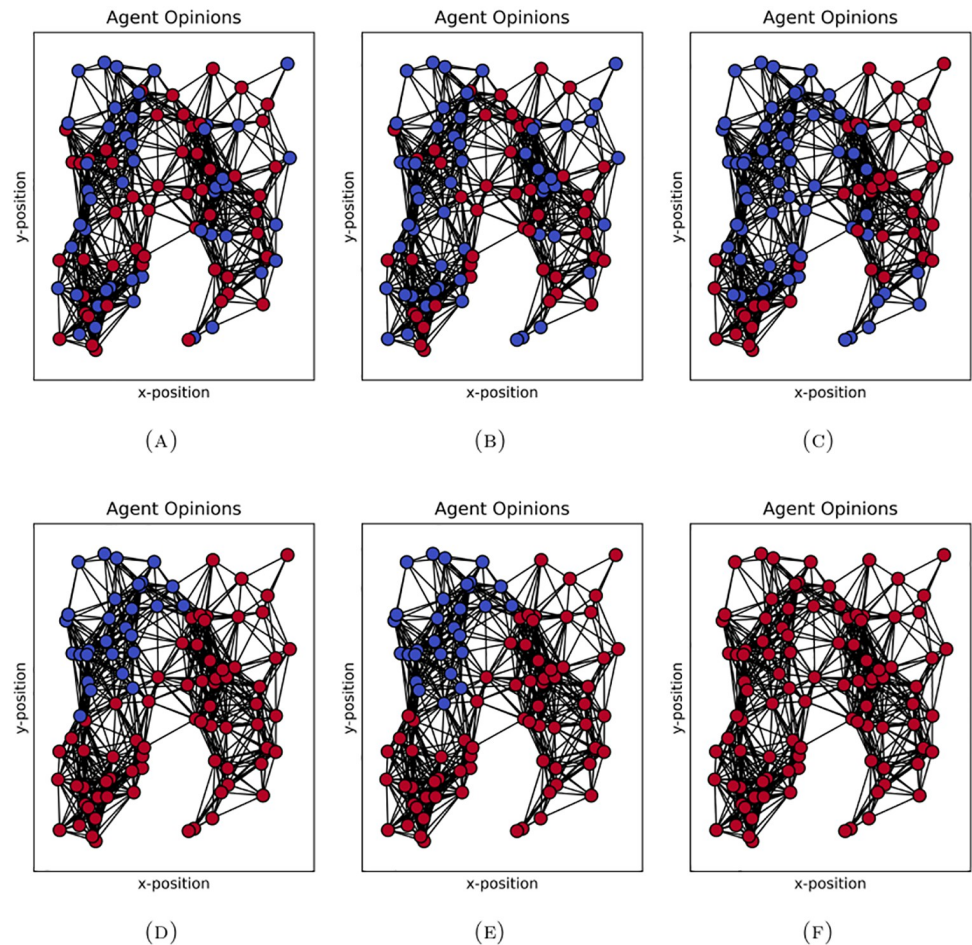


Fig 4. Opinions in simulation run with metastable behaviour at timestamp (A) $t = 1$, (B) $t = 100$, (C) $t = 10^3$, (D) $t = 10^4$, (E) $t = 10^6$, and (F) $t \approx 5.82 \times 10^6$. Note the group with opinion $o = -1$ (blue) forms around $t = 10^4$ and switches to $o = 1$ (red) after $t = 10^6$. The corresponding total number of agents holding opinion $o = 1$ is plotted in Fig 3. In this simulation run $r_g = 0.25$, the other parameters are as in §4.1.

<https://doi.org/10.1371/journal.pone.0313951.g004>

2.2 Relationship to voter model

It is not clear from the simulations presented by BO or the simulations we have executed that consensus occurs with probability one. Indeed, polarisation may seem persistent because many simulation runs ended in a state of polarisation in both sets of numerical simulations. We know that consensus will be reached asymptotically, but how long the process may be in a state of polarisation is not addressed by Theorem 1. To explore the stability of polarisation, we employ a separation of time scales argument which relates the ARLOD model to a different Markov chain, namely, the voter model.

It is well known [40–45] that reinforcement learning dynamics can be described by the replicator dynamics in the continuous time limit, using a separation of times scales between agent learning and strategy adjustment. We now present a similar relationship between the ARLOD model and the jump chain (discrete time version) of the voter model [21] on a finite topology and in the case of two opinions. It is known that the voter model on scale-free

networks [33, 34], and small-world networks [35, 36] exhibits metastable polarisation and stable consensus.

2.2.1 Discrete time voter model. In the voter model, nodes on a graph have an opinion, which may take one of two values $-1, 1$. Repeatedly, a node is selected at random from the set of all nodes. This node performs an update in which it selects one of its neighbours and copies whichever opinion they have. Time may be indexed by each such round, or by a collection of rounds in which on average each node is selected once (on the order of the population size). The version we discuss uses the former indexation of time.

We define the discrete time voter model as a Markov chain $(X_t)_{t \geq 0}$ with $t \in \mathbb{N}$. As such, we define the graph on which the voter model is to take place $G = (V, E)$, with V the set of vertices (voters) and E the set of edges (connections between voters). The number of voters is $|V| = N$ and we endow each vertex i with an opinion $o_i \in \{-1, 1\}$ for $i \in \{1, \dots, N\}$. As a result, the state space of the system is all possible assignments of each vertex to an opinion: $\mathcal{S} := \{-1, 1\}^N$.

We denote the unit vector of length N with a one at the l -th entry and zeros everywhere else, as \mathbf{e}_l for $l \in \{1, 2, \dots, N\}$. The transition probability from state $\eta \in \mathcal{S}$ to state $\zeta \in \mathcal{S}$ is denoted $P_{\eta, \zeta} := \mathbb{P}(X_{t+1} = \zeta \mid X_t = \eta)$ and is given by

$$P_{\eta, \zeta} = \begin{cases} 0 & \text{if } \|\eta - \zeta\|_1 > 2, \\ \frac{1}{N} \left(\frac{1}{2} - \frac{o_l}{2d_l} \sum_{k \in N(l)} o_k \right) & \text{if } \|\eta - \zeta\|_1 = 2, \text{ and } \zeta = \eta - 2o_l \mathbf{e}_l, \\ 1 - \sum_{\zeta \neq \eta} P_{\eta, \zeta} & \text{if } \zeta = \eta. \end{cases} \tag{8}$$

Here d_l is the degree of voter $l \in V$ and $N(l) = \{u : (u, l) \in E\}$ is their neighbourhood in the graph G .

Informally, the transition probability in (8) is simply the uniform probability of agent $i \in \{1, \dots, N\}$ being chosen, multiplied by the probability of them selecting a neighbour (uniformly at random) holding opinion $-o_j$. All transitions from $\eta \in \mathcal{S}$ to $\zeta \in \mathcal{S}$ in which the two states η and ζ differ in more than one position occur with probability zero.

Given a starting assignment of opinions to voters $\eta \in \mathcal{S}$, the voter model is the Markov process $(X_t)_{t \geq 0}$ that is Markov (δ_η, P) , taking values in \mathcal{S} . Here δ_η is the delta function. Alternatively, given a distribution of the possible starting assignments of opinions to voters λ such that $\mathbb{P}(X_0 = \eta) = \lambda_\eta$ for each $\eta \in \mathcal{S}$, the voter model is Markov (λ, P) .

The dynamics of the voter model are illustrated in Fig 5. In this example, we consider 5 voters, $V = \{1, \dots, 5\}$ with connections $E = \{(1, 2), (1, 3), (1, 4), (2, 3), (2, 5), (3, 5), (4, 5)\}$ and initial opinions $X_0 = [-1, 1, 1, -1, 1]$. We show the transitions conditioned on voter 1 being selected to copy the opinion of one of their neighbours. In particular, if voter 1 selects voters 2 or 3 they switch their opinion and if they select voter 4 they keep their current opinion. These transitions occur with probability $2/3$ and $1/3$, respectively.

2.2.2 ARLOD model in batches. The concept of multi-agent learning in batches has been explored in its own right [46–48]. It may be interpreted as a separation of time scales. That is, the rate at which agents learn about the behaviour of the environment or the other agents is faster than the rate at which they adjust their behaviour. Practically, it may be implemented by defining a batch size $b \in \mathbb{N}$ which constitutes a number of rounds in which the agent keeps their behaviour fixed and collects samples from their environment. At the end of this batch, the belief of the agent is updated using all the observations made during the batch.

To establish the link between the ARLOD model (with sophisticated agents) from computational sociology and the voter model (from sociophysics), we define the preference vector at




Fig 5. Illustration of the voter model dynamics. We show the transition probabilities conditioned on voter 1 being selected to copy the opinion of one of their neighbours at random.

<https://doi.org/10.1371/journal.pone.0313951.g005>

time $t \in \mathbb{N}$: Y_t whose elements are:

$$Y_t(i) = \begin{cases} 1 & \text{if } Q_1^i(t) \geq Q_{-1}^i(t), \\ -1 & \text{otherwise.} \end{cases} \tag{9}$$

It takes values in the state space $\mathcal{S} := \{-1, 1\}^N$.

Note that we use the weak inequality in (9), though in the limit of interest, equality occurs with probability zero. We define the preference vector for the batched model as $(Y_t^b)_{t \geq 0}$. In essence, the batched model is a biased realisation of ARLOD; in the batch at time t an agent is chosen to express their opinion to a neighbour as often ($b_t \in \mathbb{N}$ times) as is needed for them to have the same opinion preference. This occurs in the ARLOD model at probability $(1/N)^{b_t}$.

Now we define the batch learning version of the ARLOD model. In particular, agent i chosen to express their opinion in batch $t \in \mathbb{N}$ will express their opinion to their chosen neighbour j in a batch of size $b_t \in \mathbb{N}$.

That is, the dynamics follow the steps:

1. At time t , an agent $i \in \{1, 2, \dots, N\}$ is selected uniformly at random from the population.
2. This agent i chooses a neighbour j from their neighbourhood $N(i)$ uniformly at random.
3. Then follow a sequence of subrounds indexed $s = 1, \dots, b_t$. Because agent i is the only agent who can adjust their belief in this batch, we denote agent i 's Q-values in the subround s by $Q^i(s)$ and their opinion preference Y_s^i (with $Q^i(0) = Q^i(t)$ and $Y_0^i = Y_t^i(i)$). In each subround, agent i expresses an opinion to agent j , following the rules of the ARLOD model:
 - expressing their preferred opinion $o_i(s) = Y_s^i$ at probability $1 - \epsilon$,
 - expressing their disfavoured opinion $o_i(s) = -Y_s^i$ at probability ϵ , and
 - incorporating agent j 's honest response $R_j(s)$ into their Q^i value $Q_{o_s}^i$.

Now we define the random batch size $b_t = \min\{s : Y_s^i = Y_t^i(j)\}$, i.e., the number of subrounds required until agent i 's preference matches that of agent j .
4. Agent i updates their Q-values: $Q^i(t + 1) \leftarrow Q^i(t + b_t)$.

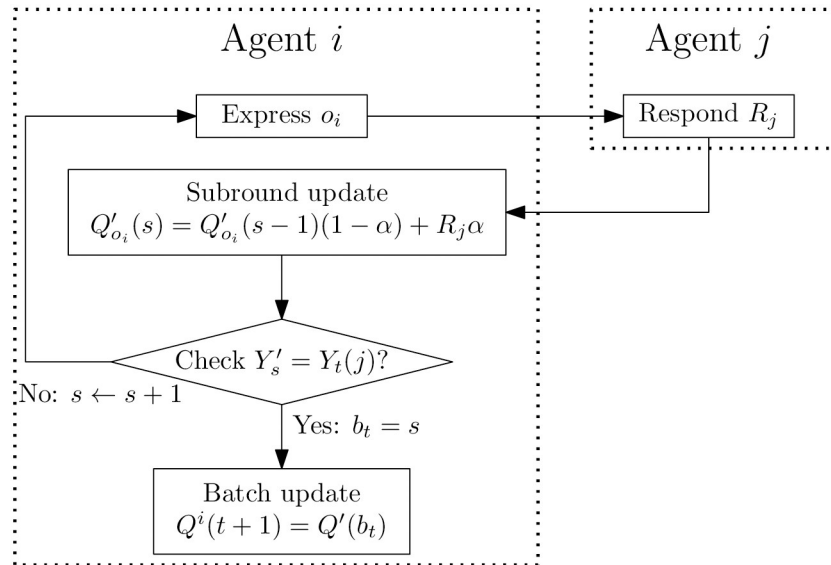


Fig 6. The dynamics in one time step of the batched version of the ARLOD model at a high level of abstraction. Agent i expresses an opinion b_t times to their neighbour agent j who responds each time. Thereafter, agent i updates their Q-values with all the feedback they received.

<https://doi.org/10.1371/journal.pone.0313951.g006>

We use this perhaps unconventional construction because the techniques in [49] are not applicable here, as the states are not lumpable.

On a high level, the procedure of one such time step is depicted in Fig 6.

2.2.3 Relationship between the ARLOD model and the voter model. We now state the main result of this section, which relates a batch learning version of the ARLOD model to the discrete time version of the voter model.

Theorem 2. For any initial assignment of Q-values resulting in preference vector $Y_0 = \eta_0 \in \mathcal{S}$, the random process tracking the change of the preference vector $(Y_t^b)_{t \geq 0}$ in the batch version of the ARLOD model on graph G converges in distribution to the voter model on the same graph:

$$\mathbb{P}(Y_t^b = \eta) = \mathbb{P}(X_t = \eta), \quad \forall \eta \in \mathcal{S}, \tag{10}$$

with $(X_t)_{t \geq 0}$ Markov (δ_{η_0}, P) with P as defined in (8).

The proof is provided in Appendix C and relies on the fact that an agent will receive enough feedback to make the ordering of their Q-values match that of their neighbour in finite time. Thus, we have shown that under a particular separation of time scales, the ARLOD model behaves like the discrete time voter model on a finite graph with two opinions. The construction of the batched ARLOD model and its relation to the voter model ensures that any state that is metastable in the voter model will also be metastable in the ARLOD model. This is because any realisation of events in the batched ARLOD model also occur with positive probability in the standard ARLOD model.

2.3 Instability of consensus and ergodicity of symmetric reinforcement learning

We now introduce a new model based closely on the ARLOD model, with a subtle difference: both agents involved in an interaction express their opinion in the same way and update their Q-values as a result of what they observe. Because now the roles of the two agents are indistinguishable, we call this the *symmetric* reinforcement learning for opinion dynamics (SRLOD) model.

2.3.1 The symmetric reinforcement learning opinion dynamics (SRLOD) model. A population of $n \in \mathbb{N}$ agents is embedded in a random (connected) geometric network topology. In each discrete time step $t \in \mathbb{N}_{\geq 0}$ an edge $(i, j) \in E$ is selected uniformly at random. The two agents on either end of this edge i and j express an opinion to one another $o_i(t), o_j(t) \in \{-1, 1\}$. Subsequently, both agents update the Q-value of their expressed opinion as follows:

$$q_{o_i(t)}^i(t + 1) = q_{o_i(t)}^i(t)(1 - \alpha) + \alpha o_i(t) o_j(t) \tag{11}$$

$$q_{o_j(t)}^j(t + 1) = q_{o_j(t)}^j(t)(1 - \alpha) + \alpha o_i(t) o_j(t), \tag{12}$$

where $\alpha \in (0, 1)$ is the learning rate. To differentiate it from the ARLOD model, we let $q_o^i(t)$ denote the Q-value agent $i \in \{1, \dots, N\}$ has for opinion $o \in \{-1, 1\}$ at time $t \in \mathbb{N}$. The Q-value of the opinion the agents did not express is not updated. We call the opinion o such that $q_o^i(t) > q_{-o}^i(t)$ agent i 's preferred opinion. We assume that both agents express their preferred opinion with probability $1 - \epsilon$ (called exploitation) and express their disfavoured opinion with probability ϵ (called exploration).

The difference thus between this model and the original model is only that instead of a one-sided interaction, both agents may explore and learn from the interaction each round.

2.3.2 Instability of consensus in the SRLOD model. We show that consensus is no longer stable in this model.

Lemma 3 (Consensus is not stable). *If there exists a time $t_0 \in \mathbb{N}$ such that $q_o^i(t_0) > q_{-o}^i(t_0)$ for some opinion $o \in \{-1, 1\}$ and each agent $i \in \{1, 2, \dots, N\}$, then*

$$\mathbb{P}(q_o^i(t) > q_{-o}^i(t), \forall t \geq t_0) = 0. \tag{13}$$

The proof of Lemma 3 is presented in Appendix D. This and the next result depend on the fact that any sequence of actions has positive probability in this model because *both* agents learn from an interaction and explore with probability $\epsilon > 0$. In particular, the probability of any finite sequence of actions of length $l < \infty$ occurs with a probability bounded from below by $p(l)$:

$$p(l) = \left(\frac{1}{|E|}\right)^l (\epsilon^2)^l > 0. \tag{14}$$

Consensus not being a stable state is a fundamental difference between the symmetric and the asymmetric model. To elucidate this difference, we introduce the preference vector, y_t , of length N , whose i -th element takes the value:

$$y_t(i) = \begin{cases} 1 & \text{if } q_1^i(t) \geq q_{-1}^i(t) \\ -1 & \text{otherwise.} \end{cases} \tag{15}$$

The preference vector describes which opinion (1 or -1) each agent $i \in \{1, 2, \dots, N\}$ favours. The dynamics of the preference vector are ergodic in the symmetric model.

Proposition 1 (Time-evolution of the preference vector is ergodic) The probability of the preference vector transitioning in finite time between any two states $\eta, \zeta \in \mathcal{S}$ is positive, i.e.,

$$\mathbb{P}(y_{t_0+k} = \zeta \mid y_{t_0} = \eta) > 0, \quad \text{for all } \eta, \zeta \in \mathcal{S}, \text{ and some } k < \infty. \quad (16)$$

To prove this, one first delineates a finite sequence between any two states $\eta, \zeta \in \mathcal{S}$ and observes that the probability of these sequences is positive by (14). For any two states there is a finite sequence of events which leads from one to the other as all agents take both actions with positive probability, and can always switch their belief in a finite number of rounds. Thus, all states communicate with one another. The ergodicity of the SRLOD model is illustrated in Appendix F. The SRLOD model being ergodic means that looking for the probability of consensus, or polarization is no longer as straightforward because both of these occur with probability one over the infinite time horizon. Instead, it is reasonable, for example, to calculate or estimate the stationary distribution of the model which gives insights into the relative time spent in consensus and polarized states. This illustrates how the tools used to study a model differ based on the seemingly innocuous assumption of asymmetry in the agent-to-agent interaction.

3 Discussion

We have analysed the ARLOD model of social learning put forth by Banisch and Olbrich [22]. Our first main theorem shows that consensus is reached asymptotically with probability one for any finite and connected population structure. In particular, this is in contrast with the persistence of polarisation originally reported for that model. A small modification of that model, based on symmetrizing the interaction-learning relation between the agents, results instead in ergodic dynamics, which thus destabilizes consensus somewhat. This result mirrors the difference between the voter model and the *noisy* voter model, in which a random probability of switching one's opinion is introduced [50, 51].

The highlighted importance of network structure in the original article [22] warrants attention. The theoretical arguments we present here to show that a) consensus is the only stable state in the original model and b) that the symmetrized model is ergodic required only that the network is connected. Thus, *qualitatively*, the assumption of network structure is not very important. We do, however, see that it plays an important role *quantitatively* in the time taken for consensus to be reached. It is likely that the metastability of polarization emerges because of strong community structure in the network. This is in accordance with previous findings of the effect of network structure on the timescales of the resulting dynamics (see for instance [52, 53]) Many studies on polarisation and other social dynamics focus on the importance of network structure. In addition to investigating the effects of networks, it may also be important to disentangle which outcomes of the model are truly caused by networks structure and which outcomes are the result of other—more implicit—modelling decisions, such as asymmetry in the agent-to-agent interaction.

Having proved that the polarisation observed in the ARLOD model is not stable and that consensus is guaranteed, we turn to the original research question. What causes stable polarisation? We provide conditions (a systematic biasing) for which the ARLOD model converges to the voter model. Polarisation can be metastable in the voter model, and, by their relation, also in the reinforcement learning model. This (relationship between the ARLOD model and the voter model) bridges multiagent learning models and models well studied in sociophysics and theoretical biology.

Our results raise questions regarding the possibility of finding a model of opinion dynamics excluding repulsive forces and allowing for stable polarisation. Can we say that a reasonable

model of opinion dynamics should exhibit stable (as defined in Definition 1) polarisation? Is the polarisation we observe around us stable or metastable? Future research is required to give an example of such a model or a proof that it does not exist. These questions might be explored by investigating learning in the ‘real world’ (to identify appropriate α and ϵ) as well as the influence of parameter values α and ϵ in the ARLOD model. It could be that ‘real world’ learning is such that consensus would be reached quickly under the ARLOD model, indicating that a more realistic model requires additional elements. Alternatively, it may be that the parameters of the ‘real world’ are such that the time it takes to exit the metastable polarised state is so long that differentiating between metastable and stable polarisation in the real world is difficult.

A limitation of the models we study is that the memory of the agents is entirely implicit because we use stateless Q-learning. Explicit inclusion of memory may be done by using Q-learning *with* states, where each state corresponds to the last action taken by each of an agents neighbours. This would complicate the analysis: If additionally to memory, an agent knows the identity of the neighbour they are expressing their opinion to, it is possible that polarization becomes absorbing (and therewith stable). Note that care would have to be taken to determine how an agent responds to an opinion to avoid decoupling the dynamics of each pair of agents from other interactions. Another limitation of the model relates to the isolation of the dynamics from other influences. Effects other than social influence that may be driving agent opinions are an internal cognitive process related to their opinion such as in [15, 54], or pressure applied by mass media to follow a certain opinion [26, 55–57]. Finally, the assumption that the agents of the model are fixed (no new agents enter, or old ones leave) can be seen as unrealistic. Important to note that changing this assumption may change the outcome of the analysis. In particular if new agents have random Q-values, this destabilizes consensus in the ARLOD model.

4 Methods

4.1 ARLOD simulation settings

We have chosen the parameter settings based on the following considerations. A greater number of agents means that more rounds are required to select each agent sufficiently often to reach consensus. On the other hand, a smaller learning rate increases what ‘sufficiently often’ means per agent, as indicated in (4). To strike a balance between these effects, we set $N = 100$ and $\alpha = 0.25$. Following BO, we set $\epsilon = 0.1$ and initialise the Q-values uniformly in $[-0.5, 0.5]$. The radius for the random geometric graph model $r_g \in [0.25, 0.5]$ is selected to exhibit a range of behaviour, focusing on connected graphs. We have chosen the maximum time to simulate (10×10^9 rounds) and the number of simulation iterations (500 iterations) to be significantly greater than those used by BO (2×10^6 rounds and 100 iterations). This allows the simulation to reach consensus more frequently, which we know occurs eventually with probability one (by Theorem 1).

4.2 Random connected geometric graphs

The algorithm to generate a connected random graph is provided in Appendix E. We use the subroutine for the generation of a random geometric network from the Python NetworkX package [58]. For a detailed discussion on random geometric graphs and their properties, the interested reader is referred to [38, 59]. The random geometric graph model is popular in the context of social dynamics because it mimics the homophily of real social networks as claimed by [60].

The general idea of the random geometric graph is to distribute the desired number of nodes randomly in Euclidean space (we use $[0, 1]^2$) and fixes a radius r_g . Subsequently, any

nodes u, v that are distance $d(u, v) < r_g$ from one another are connected by an edge (u, v) . Because we are interested in connected networks, we simply repeat the standard procedure until a connected graph is sampled. We take care to only use r_g for which the probability of sampling a connected graph is sufficiently high (as described in §4.1).

Proof of Lemma 1

We proceed by induction on time. Suppose that $Q_o^i(t_0) > Q_{-o}^i(t_0)$ at time $t_0 \geq 0$, for some opinion $o \in \{-1, 1\}$, and each agent $i \in \{1, 2, \dots, N\}$.

The base case is that in round $t_0 + 1$, the ordering of all the Q-values will remain the same.

In round t_0 , any agent $i \in \{1, 2, \dots, N\}$ may be chosen to express their opinion to one of their neighbours.

Case 1. Suppose they exploit their preferred opinion (the one with greater Q-value). Any agent they express their opinion to, has the same ordering among their Q-values by the conditions of the lemma, and so responds with an action that leads to a positive reward. Thus,

$$Q_o^i(t_0 + 1) \geq Q_o^i(t_0), \tag{17}$$

the Q-value of the preferred opinion in round $t_0 + 1$ is at least as great as in t_0 .

Case 2. Suppose they explore by taking the action with lesser Q-value. Any neighbour they express this opinion to responds honestly. By the assumption, all agents have the same Q-value ordering, so the honest response to exploration is an action that leads to a punishment. Thus,

$$Q_{-o}^i(t_0 + 1) \leq Q_{-o}^i(t_0), \tag{18}$$

the Q-value of the disfavoured action in round $t_0 + 1$ is lower than or equal to what it was in round t_0 . This is true because the Q-values are initialised to be in $[-1, 1]$ and will stay therein indefinitely by the updating prescribed.

This proves the base case (as this holds for all agents that could have been chosen in round t_0): $Q_o^i(t_0 + 1) > Q_{-o}^i(t_0 + 1)$ for all agents $i \in \{1, 2, \dots, N\}$.

In the induction step we assume it is true until rounds $t_0 + n$ for $n > 0$. To show that it is true for all rounds up until $t_0 + n + 1$, we simply follow the same procedure as in the base case but for the game in round $t_0 + n$ which determines the Q-values in round $t_0 + n + 1$.

Proof of Lemma 2

First, we delineate a sequence of events of finite length which may lead from any state to consensus. Secondly, we will show that this sequence of events has positive probability.

Suppose agent i favours opinion o and has a neighbour j who prefers opinion $-o$, all at time t_0 . If agent i is drawn to express their opinion to agent j every round for $L \in \mathbb{N}$ rounds and always exploits their preferred opinion, the Q-value for this opinion is given by:

$$Q_o^i(t_0 + l) = Q_o^i(t_0 + l - 1)(1 - \alpha) - \alpha, \tag{19}$$

for all $l = 1, 2, \dots, L$. A term by term comparison shows that this is bounded from above by

$$Q_o^i(t_0 + l) \leq Q_o^i(t_0 + l - 1)(1 - \alpha), \tag{20}$$

since $\alpha \in (0, 1)$. Thus, an upper bound of the Q-value in round $t_0 + l$ is given by $Q_o^i(t_0)(1 - \alpha)^l$ for all $l = 1, 2, \dots, L$ as long as $Q_o^i(t_0) > 0$. When both Q-values have the same sign, only one of them needs to be adjusted in the way described here until it changes sign.

Subsequently, if agent i is drawn to express their opinion to agent j another $M \in \mathbb{N}$ times and explores their disfavoured action in each of these rounds, this opinion's Q-value follows:

$$Q_{-o}^i(t_0 + L + m) = Q_{-o}^i(t_0 + L + m - 1)(1 - \alpha) + \alpha, \tag{21}$$

for all $m = 1, 2, \dots, M$. A term by term comparison shows that this is bounded from below by

$$Q_{-o}^i(t_0 + L + m) \geq Q_{-o}^i(t_0 + L + m - 1)(1 - \alpha). \tag{22}$$

Again a lower bound to this Q-value in round $t + 0 + L + m$ is given by $Q_{-o}^i(t_0 + L)(1 - \alpha)^m$ as long as $Q_{-o}^i(t_0 + L) < 0$.

We bound from above the number of rounds needed for any agent's opinion to be switched, by the number of rounds needed should they start as far away from one another as possible, $\mathbf{Q} = (-1, 1)$, or $(1, -1)$ and be set to cross at zero. The Q-value of the originally preferred opinion reaches $\xi \in (0, \alpha)$ at least by the lowest integer r which satisfies:

$$Q_o^i(t_0)(1 - \alpha)^r \leq \xi, \tag{23}$$

if they express this preferred opinion in each round. Dividing by $Q_o^i(t_0)$, taking the logarithm on both sides and rearranging we get

$$r = \left\lceil \frac{\log(\xi/Q_o^i(t_0))}{\log(1 - \alpha)} \right\rceil. \tag{24}$$

By a similar procedure we see that the Q-value of the originally disfavoured opinion reaches $-\epsilon$ after a further r interactions (of exploring in each subsequent round). After two more interactions in which the agent expresses each opinion once, the Q-value ordering has switched:

$$Q_{-o}^i(t_0 + 2r + 2) \geq -\xi(1 - \alpha) + \alpha > \xi(1 - \alpha) - \alpha \geq Q_o^i(t_0 + 2r + 2), \tag{25}$$

as long as $\alpha > \xi$, which is satisfied by an appropriate choice of ξ .

The number of rounds this takes is $2r + 2$. The probability that this happens is bounded by the probability of the agent j being drawn to express their opinion to agent k $2r + 2$ times, multiplied by the probability that they take the required action in each round. This is a lower bound because it does not matter whether agent i first exploits $r + 1$ times and then explores $r + 1$ times in that order. It only matters that there is a total of $r + 1$ explorations and exploitations in the $2r + 2$ rounds. Thus, the probability p_{switch} of one agent switching their opinion (if they have at least one neighbour that disagrees with them) is lower bounded by

$$p_{\text{switch}} > (1 - \epsilon)^{r+1} \epsilon^{r+1} \left(\frac{1}{N(N - 1)} \right)^{2r+2} > 0. \tag{26}$$

Here, agent i is drawn to express their opinion with probability $1/N$ and we bound the probability that they express this opinion to agent k from below by $1/(N-1)$ as that is the maximum possible degree for any agent in the network. This probability is greater than zero simply because it is a finite product of positive numbers.

In a connected population of N agents which is not yet in consensus, there is always at least one edge which has an agent who prefers opinion o on one side and opinion $-o$ on the other side. Furthermore, in the initial state there are at most $N-1$ agents who prefer the 'wrong' opinion at time t_0 . So with probability $p > p_{\text{switch}}^{N-1} > 0$, in $(N - 1)(2r + 2) < \infty$ rounds all N agents hold the same opinion.

C Proof of Theorem 2

At time $t = 0$, by definition, we have that $\mathbb{P}(Y_0 = \eta_0) = 1$ which is also true for X_t under δ_{η_0} : $\mathbb{P}(X_0 = \eta_0) = 1$.

Next we show that for $t \in \mathbb{N}$,

$$\mathbb{P}(Y_{t+1} = \eta \mid Y_t = \eta_t, \dots, Y_0 = \eta_0) = \mathbb{P}(Y_{t+1} = \eta \mid Y_t = \eta_t) = P_{\eta_t, \eta}. \tag{27}$$

The Y_{t+1} 's independence on Y_{t-1}, \dots, Y_0 follows from the fact that in the batch at time $t \in \mathbb{N}$ the agents determine their behaviour entirely from the state Y_t . Expressing agents express their favoured opinion with probability $1-\epsilon$ and express their disfavoured opinion with probability ϵ . Responding agents always do so honestly, rewarding their favoured opinion and punishing their disfavoured opinion.

Note that $\mathbb{P}(Y_{t+1} = \eta \mid Y_t = \eta_t) = 0$ whenever $\|\eta - \eta_t\|_1 > 2$, just as in (8). This is because as soon as $\|\eta - \eta_t\|_1 > 2$ we have that more than one agent has switched their opinion after the batch at time t . This is impossible because only one agent updates their Q-values during a batch.

We proceed in two cases, one when $\|\eta - \eta_t\|_1 = 2$ and the other when $\|\eta - \eta_t\|_1 = 0$. Note that $\|\eta - \eta_t\|_1 \neq 1$ because $\eta, \eta_t \in \mathcal{S} = \{-1, 1\}^N$.

Case 1. $\|\eta - \eta_t\|_1 = 2$: In this case states η and η_t differ in exactly one position which, without loss of generality, we label l . In order for agent l to switch their favoured opinion in batch t , they must be selected uniformly at random to express their opinion to one of their neighbours. This happens at probability $1/N$.

If an agent expresses their opinion to an agent who favours opinion $-o_b$, agent l , gets punished for each time they express their favoured opinion and rewarded for each time they express their disfavoured opinion. Thus, if agent l expresses their favoured opinion at least $r + 1$ times and their disfavoured opinion at least $r + 1$ times, where

$$r = \left\lceil \frac{\log(\epsilon)}{\log(1 - \alpha)} \right\rceil, \tag{28}$$

then their opinion will have switched (see the proof of Lemma 2). Thus, we can lower bound the probability of the agent switching their opinion after $2r + 2$ rounds by

$$p_{\text{switch}} \geq \epsilon^{r+1}(1 - \epsilon)^{r+1}. \tag{29}$$

Then the probability that the agent does not switch their opinion in finite time is upper bounded by $\lim_{k \rightarrow \infty} (1 - p_{\text{switch}})^k = 0$. This means that if the agent has selected to express their opinion to an agent who favours opinion $-o_b$, they will switch their opinion, and this will happen in finite time.

The probability of agent l switching their opinion is given by the probability that they select a neighbour favouring the opposite opinion to theirs. As before, we denote d_l as the degree of agent l . Furthermore, we denote with $a(l)$ and $c(l)$ the number of agents in l 's neighbourhood who are in agreement and contradiction with l respectively. Then, because agents select a neighbour uniformly at random, the probability of l switching their opinion is

$$\mathbb{P}(o_l \rightarrow -o_l) = \frac{c(l)}{d_l}. \tag{30}$$

This may be rewritten and rearranged as follows:

$$\mathbb{P}(o_l \rightarrow -o_l) = \frac{c(l)}{2d_l} + \frac{c(l)}{2d_l} \tag{31}$$

$$= \frac{c(l)}{2d_l} + \frac{d_l - a(l)}{2d_l} \tag{32}$$

$$= \frac{1}{2} \left(1 + \frac{c(l) - a(l)}{d_l} \right). \tag{33}$$

To extract this from Y_t , notice that it can be rewritten as

$$\mathbb{P}(o_l \rightarrow -o_l) = \frac{1}{2} \left(1 - \frac{Y_t(l)}{d_l} \sum_{k \in N(l)} Y_t(k) \right), \tag{34}$$

where $N(l)$ is the neighbourhood of agent l in the graph G . Multiplying this by $1/N$, the probability that agent l is selected to express their opinion in the first place, we get precisely P_{η_t} , η as required.

Case 2. $\|\eta - \eta_t\|_1 = 0$: In this case, an agent is selected to express their opinion, and they do so to a neighbour who is in agreement with them (in which case $b_t = 1$). Given that agent l is selected, this happens with probability

$$\mathbb{P}(o_l \rightarrow o_l) = 1 - \frac{1}{2} \left(1 - \frac{Y_t(l)}{d_l} \sum_{k \in N(l)} Y_t(k) \right). \tag{35}$$

This is exactly $1 - NP_{\eta_t}$, η for $\eta = \eta_t - 2Y_t(l)\mathbf{e}_l$. Summing over all agents that might be selected and multiplying by the probability of selecting those agents, we get the required probability of transitioning to the same state, $P_{\eta_t, \eta_t} = 1 - \sum_{\eta \neq \eta_t} P_{\eta_t, \eta}$.

D Proof of Lemma 3

Proof. We show that a sequence of events which leads from consensus to not-consensus is of finite length and positive probability.

Observe that each agent explores (expresses and reinforces the disfavoured opinion) at probability $\epsilon > 0$. Observe also that if this disfavoured opinion is reinforced maximally $\kappa + 1$ times with

$$\kappa = \left\lceil \frac{\log(\epsilon)}{\log(1 - \alpha)} \right\rceil, \tag{36}$$

which is finite, and similarly if their preferred opinion is punished $\kappa + 1$ times, that they switch ordering of opinion Q-values. In similar fashion to as it was shown in the proof of Lemma 2 for the asymmetric model.

Given a sequence of agent actions, the probability that they take the action in some round required by the sequence, is always bounded from below by $\epsilon > 0$. This is because they express their disfavoured opinion at probability ϵ and their favoured opinion at probability $1 - \epsilon > \epsilon$.

This means that the probability of any finite sequence of actions of length $l < \infty$ occurs at a probability bounded from below by $p(l)$ defined in (14). Thus, the probability of an agent exploring and being rewarded $\kappa + 1$ times and exploiting and being punished $\kappa + 1$ times is positive because this is a sequence of events with length $2\kappa + 2$. This is the maximal length

sequence which leads to one agent changing the ordering of their Q-values. So we have that the probability that such a switch *never* happens (the probability of consensus for all $t \geq t_0$):

$$\mathbb{P}(q_o^i(t) > q_{-o}^i(t), \forall t \geq t_0) \leq \lim_{k \rightarrow \infty} (1 - p(2k + 2))^k = 0. \quad (37)$$

Therefore consensus is not absorbing (and not a stable state).

E Algorithm to generate a geometric random graph

Algorithm 1: Generate connected random geometric graph with N nodes and radius r_g .

Input: N : number of nodes, r_g : radius

Output: G : connected random geometric graph

```

1 Check ← True;
2 while Check do
3   G ← empty graph;
4   for i ← 1 to N do
5     Add node i with random coordinates in  $[0, 1]^2$  to G;
6   end
7   for each pair of nodes (u, v) in G do
8     if distance(u, v)  $\leq r_g$  then
9       Add edge (u, v) to G;
10    end
11  end
12  if G is connected then
13    Check ← False
14  end
15 end

```

F Illustration of the ergodicity in the SRLOD model

The dynamics of the SRLOD model are ergodic by Proposition 1, meaning that the process can reach all states from all other states. In the ergodic setting on a finite state space, one can look for a stationary distribution. That is a probability distribution over the states reporting the probability of observing each state as $t \rightarrow \infty$. Thus, the questions one might ask of the model changes from ‘What is the probability of consensus?’ to ‘What proportion of time does the system spend in each state?’ To illustrate the ergodicity of the SRLOD model, we show a simulation run in which consensus was reached on both opinions in Figs 7 and 8.




Fig 7. The number of agents holding opinion $o = 1$ in a simulation run of the SRLOD model plotted with time on a logarithmic scale. Notice the switch from consensus on opinion $o = 1$ to opinion $o = -1$ shortly after $t = 10^6$.

<https://doi.org/10.1371/journal.pone.0313951.g007>

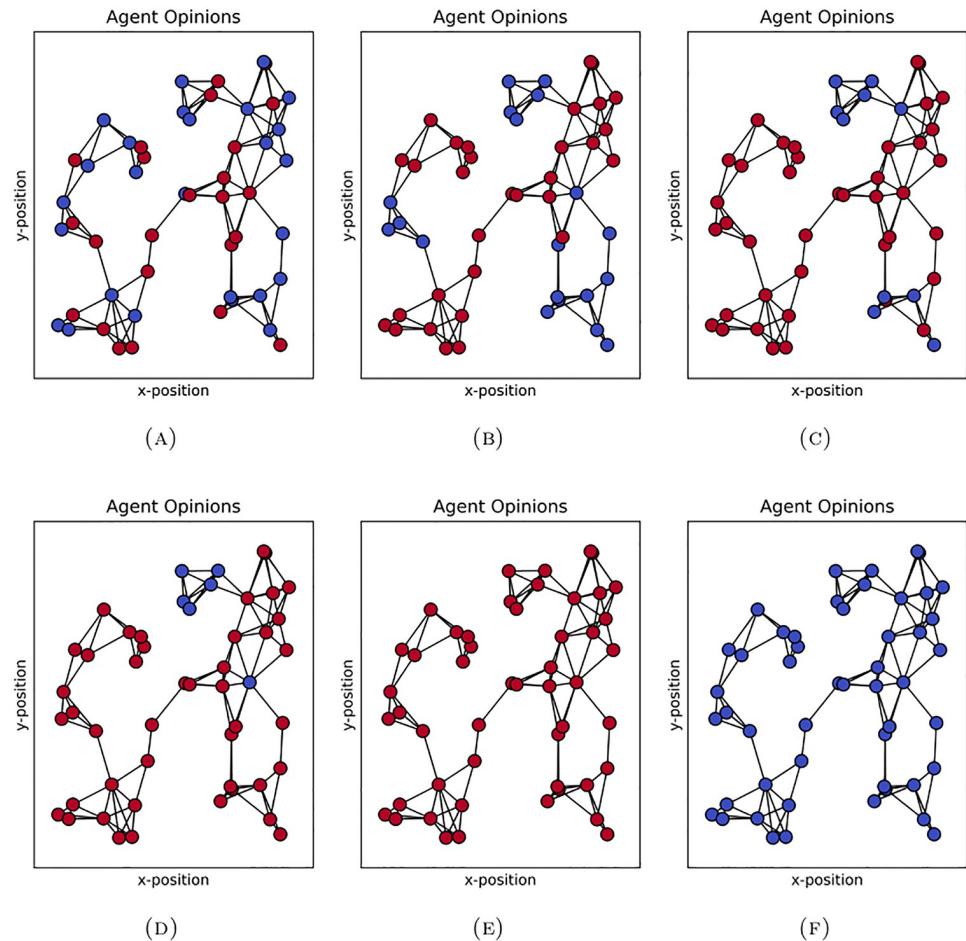


Fig 8. Opinions in simulation run of the SRLOD model at rounds (A) $t = 1$, (B) $t = 10^3$, (C) $t = 10^4$, (D) $t = 2 \times 10^5$, (E) $t = 3 \times 10^5$, and (F) $t = 2 \times 10^6$. Notice the switch from consensus on opinion $o = 1$ (red) to opinion $o = -1$ (blue) between $t = 3 \times 10^5$ and $t = 2 \times 10^6$.

<https://doi.org/10.1371/journal.pone.0313951.g008>

In this simulation we set $\alpha = 0.25$, $\epsilon = 0.2$, $N = 50$ and $r_g = 0.2$. Because consensus is no longer stable, we stop the simulation manually at $t = 8 \times 10^6$. The number of agents holding opinion $o = 1$ is plotted on a logarithmic timescale in Fig 7 while the state of the network is shown for telling rounds in Fig 8. The reason for reducing the number of agents in this simulation is to speed up the dynamics thus making it possible to observe the phenomena in a relatively short simulation.

Author Contributions

Conceptualization: Benedikt V. Meylahn.

Formal analysis: Benedikt V. Meylahn.

Investigation: Janusz M. Meylahn.

Methodology: Benedikt V. Meylahn, Janusz M. Meylahn.

Visualization: Benedikt V. Meylahn.

Writing – original draft: Benedikt V. Meylahn.

Writing – review & editing: Janusz M. Meylahn.

References

1. Abelson RP. Mathematical models of the distribution of attitudes under controversy. *Contributions to Mathematical Psychology*. 1964; 14:1–160.
2. Deffuant G, Neau D, Amblard F, Weisbuch G. Mixing beliefs among interacting agents. *Advances in Complex Systems*. 2000; 3:87–98. <https://doi.org/10.1142/S0219525900000078>
3. Hegselmann R, Krause U. Opinion dynamics and bounded confidence models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*. 2000; 5(3).
4. Weisbuch G. Bounded confidence and social networks. *The European Physics Journal B*. 2004; 38:339–343. <https://doi.org/10.1140/epjb/e2004-00126-9>
5. Gómez-Serrano J, Graham C, Le Boudec JY. The bounded confidence model of opinion dynamics. *Mathematical Models and Methods in Applied Sciences*. 2012; 22(2).
6. Flache A, Macy MW. Small worlds and cultural polarization. *The Journal of Mathematical Sociology*. 2011; 35(1-3):146–176. <https://doi.org/10.1080/0022250X.2010.532261>
7. Sobkowicz P. Discrete Model of Opinion Changes Using Knowledge and Emotions as Control Variables. *PLOS ONE*. 2012; 7(9):1–16. <https://doi.org/10.1371/journal.pone.0044489> PMID: 22984516
8. Altafini C. Consensus Problems on Networks With Antagonistic Interactions. *IEEE Transactions on Automatic Control*. 2013; 58(4):935–946. <https://doi.org/10.1109/TAC.2012.2224251>
9. Chan KMD, Duivendoorn R, Flache A, Mandjes M. A relative approach to opinion formation. *The Journal of Mathematical Sociology*. 2022; 48(1):1–41. <https://doi.org/10.1080/0022250X.2022.2036142>
10. Burke MF, Searle C. Quantitatively modelling opinion dynamics during elections. *ORION*. 2022; 38(2):123–146.
11. Flache A, Torenlvlied R. When will they ever make up their minds? The social structure of unstable decision making. *The Journal of Mathematical Sociology*. 2004; 28(3):171–196. <https://doi.org/10.1080/00222500490480257>
12. Galam S. Collective beliefs versus individual inflexibility: The unavoidable biases of a public debate. *Physica A: Statistical Mechanics and its Applications*. 2011; 390(17):3036–3054. <https://doi.org/10.1016/j.physa.2011.03.021>
13. Yildiz E, Ozdaglar A, Acemoglu D, Saberi A, Scaglione A. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*. 2013; 1(4):19:1–19:30. <https://doi.org/10.1145/2538508>
14. Gaisbauer F, Olbrich E, Banisch S. Dynamics of opinion expression. *Physical Review E*. 2020; 102:042303. <https://doi.org/10.1103/PhysRevE.102.042303> PMID: 33212677
15. Meylahn BV, Searle C. Opinion dynamics beyond social influence. *Network Science*. 2024;1–27.
16. French JRP. A formal theory of social power. *Psychological Review*. 1956; 63(3):181–194. <https://doi.org/10.1037/h0046123> PMID: 13323174
17. Harary F. A criterion for unanimity in French's theory of social power. In: Cartwright D, editor. *Studies in social power*. Ann Arbor (MI): Institute for Social Research; 1959. p. 168–182.
18. DeGroot MH. Reaching a consensus. *Journal of the American Statistical Association*. 1974; 69(345):118–121. <https://doi.org/10.1080/01621459.1974.10480137>
19. Chatterjee S, Seneta E. Towards consensus: Some convergence theorems on repeated averaging. *Journal of Applied Probability*. 1977; 14(1):89–97. <https://doi.org/10.2307/3213262>
20. Clifford P, W SA. A model for spatial conflict. *Biometrika*. 1973; 60:581–588. <https://doi.org/10.1093/biomet/60.3.581>
21. Holley RA, Liggett TM. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*. 1975; 3(4):643–663. <https://doi.org/10.1214/aop/1176996306>
22. Banisch S, Olbrich E. Opinion polarization by learning from social feedback. *The Journal of Mathematical Sociology*. 2019; 43(2):76–103. <https://doi.org/10.1080/0022250X.2018.1517761>
23. Castellano C, Fortunato S, Loreto V. Statistical physics models of social dynamics. *Review of Modern Physics*. 2009; 81(2):591–646. <https://doi.org/10.1103/RevModPhys.81.591>
24. Flache A, Mäs M, Feliciani T, Chattoe-Brown E, Deffuant G, Huet S, et al. Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*. 2017; 20(4). <https://doi.org/10.18564/jasss.3521>

25. Törnberg P, Andersson C, Lindgren K, Banisch S. Modeling the emergence of affective polarization in the social media society. *PLOS ONE*. 2021; 16(10):1–17. <https://doi.org/10.1371/journal.pone.0258259> PMID: 34634056
26. Törnberg P. How digital media drive affective polarization through partisan sorting. *Proceedings of the National Academy of Sciences*. 2022; 119(42):e2207159119. <https://doi.org/10.1073/pnas.2207159119> PMID: 36215484
27. Yu C, Tan G, Lv H, Wang Z, Meng J, Hao J, et al. Modelling adaptive learning behaviours for consensus formation in human societies. *Scientific Reports*. 2016; 6(1). <https://doi.org/10.1038/srep27626> PMID: 27282089
28. Chen T, Li Q, Fu P, Yang J, Xu C, Cong G, et al. Public Opinion Polarization by Individual Revenue from the Social Preference Theory. *International Journal of Environmental Research and Public Health*. 2020; 17(3). <https://doi.org/10.3390/ijerph17030946> PMID: 32033012
29. Lorenz J, Neumann M, Schröder T. Individual attitude change and societal dynamics: Computational experiments with psychological theories. *Psychological Review*. 2021; 128(4):623–642. <https://doi.org/10.1037/rev0000291> PMID: 34060889
30. Botte N, Ryckebusch J, Rocha LEC. Clustering and stubbornness regulate the formation of echo chambers in personalised opinion dynamics. *Physica A: Statistical Mechanics and its Applications*. 2022; 599:127423. <https://doi.org/10.1016/j.physa.2022.127423>
31. Lefebvre G, Deroy O, Bahrami B. The roots of polarization in the individual reward system. *Proceedings of the Royal Society B: Biological Sciences*. 2024; 291:20232011. <https://doi.org/10.1098/rspb.2023.2011> PMID: 38412967
32. Cox JT. Coalescing random walks and voter model consensus times on the torus in Zd. *The Annals of Probability*. 1989; 17(4):1333–1366. <https://doi.org/10.1214/aop/1176991158>
33. Suchecki K, Eguíluz VM, San Miguel M. Voter model dynamics in complex networks: Role of dimensionality, disorder, and degree distribution. *Physical Review E*. 2005; 72:036132. <https://doi.org/10.1103/PhysRevE.72.036132> PMID: 16241540
34. Suchecki K, Eguíluz VM, San Miguel M. Conservation laws for the voter model in complex networks. *Europhysics Letters*. 2005; 62:228–234. <https://doi.org/10.1209/epl/i2004-10329-8>
35. Castellano C, Vilone D, Vespignani A. Incomplete ordering of the voter model on small-world networks. *Europhysics Letters*. 2003; 63(1):153. <https://doi.org/10.1209/epl/i2003-00490-0>
36. Vilone D, Castellano C. Solution of voter model dynamics on annealed small-world networks. *Physical Review E*. 2004; 69:016109. <https://doi.org/10.1103/PhysRevE.69.016109> PMID: 14995669
37. Banisch S, Shamon H. Biased Processing and Opinion Polarization: Experimental Refinement of Argument Communication Theory in the Context of the Energy Debate. *Sociological Methods & Research*. 2023; p. 00491241231186658. <https://doi.org/10.1177/00491241231186658>
38. Dall J, Christensen M. Random geometric graphs. *Physical Review E*. 2002; 66:016121. <https://doi.org/10.1103/PhysRevE.66.016121> PMID: 12241440
39. Meylahn BV, den Boer AV, Mandjes MRH. Interpersonal trust: Asymptotic analysis of a stochastic coordination game with multi-agent learning. *Chaos*. 2024; 34(6):063119. <https://doi.org/10.1063/5.0205136> PMID: 38848273
40. Phansalkar VV, Sastry PS, Thathachar MAL. Absolutely expedient algorithms for learning Nash equilibria. *Proceedings of the Indian Academy of Sciences: Mathematical Sciences*. 1994; 104(1):279–194.
41. Sastry PS, Phansalkar VV, Thathachar MAL. Decentralized learning of Nash equilibria in multi-person stochastic games with incomplete information. *IEEE Transactions on Systems, Man, and Cybernetics*. 1994; 24(5). <https://doi.org/10.1109/21.293490>
42. Börgers T, Sarin R. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*. 1997; 77(1):1–14. <https://doi.org/10.1006/jeth.1997.2319>
43. Sato Y, Akiyama E, Farmer JD. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*. 2002; 99(7):4748–4751. <https://doi.org/10.1073/pnas.032086299> PMID: 11930020
44. Sato Y, Crutchfield JP. Coupled replicator equations for the dynamics of learning in multiagent systems. *Physical Review E*. 2003; 67:015206. <https://doi.org/10.1103/PhysRevE.67.015206> PMID: 12636552
45. Sato Y, Akiyama E, Crutchfield JP. Stability and diversity in collective adaptation. *Physica D: Nonlinear Phenomena*. 2005; 210(1):21–57. <https://doi.org/10.1016/j.physd.2005.06.031>
46. Barfuss W, Donges JF, Kurths J. Deterministic limit of temporal difference reinforcement learning for stochastic games. *Physical Review E*. 2019; 99:043305. <https://doi.org/10.1103/PhysRevE.99.043305> PMID: 31108579

47. Barfuss W, Meylahn JM. Intrinsic fluctuations of reinforcement learning promote cooperation. *Scientific Reports*. 2023; 13(1). <https://doi.org/10.1038/s41598-023-27672-7> PMID: 36693872
48. Meylahn JM, Janssen L. Limiting dynamics for Q-Learning with memory one in symmetric two-player, two-action games. *Complexity*. 2022; 2022(4830491). <https://doi.org/10.1155/2022/4830491>
49. Banisch S, Lima R, Araújo T. Agent based models and opinion dynamics as Markov chains. *Social Networks*. 2012; 34(4):549–561. <https://doi.org/10.1016/j.socnet.2012.06.001>
50. Granovsky BL, Madras N. The noisy voter model. *Stochastic Processes and their Applications*. 1995; 55(1):23–43. [https://doi.org/10.1016/0304-4149\(94\)00035-R](https://doi.org/10.1016/0304-4149(94)00035-R)
51. Carro A, Toral R, San Miguel M. The noisy voter model on complex networks. *Scientific Reports*. 2016; 6(1):24775. <https://doi.org/10.1038/srep24775> PMID: 27094773
52. Centola D, Baronchelli A. The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*. 2015; 112(7):1989–1994. <https://doi.org/10.1073/pnas.1418838112> PMID: 25646462
53. Zarei F, Gandica Y, Rocha LEC. Bursts of communication increase opinion diversity in the temporal Deffuant model. *Scientific Reports*. 2024; 14:2222. <https://doi.org/10.1038/s41598-024-52458-w> PMID: 38278824
54. Giardini F, Vilone D, Conte R. Consensus emerging from the bottom-up: The role of cognitive variables in opinion dynamics. *Frontiers in Physics*. 2015; 3. <https://doi.org/10.3389/fphy.2015.00064>
55. Hoffman LH, Glynn CJ, Hugu ME, Sietman RB, Thomson T. The Role of Communication in Public Opinion Processes: Understanding the Impacts of Intrapersonal, Media, and Social Filters. *International Journal of Public Opinion Research*. 2007; 19(3):287–312. <https://doi.org/10.1093/ijpor/edm014>
56. Pansanella V, Sirbu A, Kertesz J, Rossetti G. Mass media impact on opinion evolution in biased digital environments: A bounded confidence model. *Scientific Reports*. 2023; 13:14600. <https://doi.org/10.1038/s41598-023-39725-y> PMID: 37670041
57. Van Santen N, Ryckebusch J, Rocha LEC. Social clustering reinforces external influence on the majority opinion model. *Physica A: Statistical Mechanics and its Applications*. 2024; 648:129929. <https://doi.org/10.1016/j.physa.2024.129929>
58. Hagberg AA, Schult DA, Swart PJ. Exploring network structure, dynamics, and function using NetworkX. In: Varoquaux G, Vaught T, Millman J, editors. *Proceedings of the 7th Python in Science Conference*. Pasadena, (CA) USA; 2008. p. 11–15.
59. Penrose M. *Random Geometric Graphs*. Oxford University Press; 2003. Available from: <https://doi.org/10.1093/acprof:oso/9780198506263.001.0001>.
60. McPherson M, Smith-Lovin L, Cook JM. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*. 2001; 27:415–444. <https://doi.org/10.1146/annurev.soc.27.1.415>