

An evidence-accumulating drift–diffusion model of competing information spread on networks

Original

An evidence-accumulating drift–diffusion model of competing information spread on networks / Corsin, Julien; Zino, Lorenzo; Ye, Mengbin. - In: CHAOS, SOLITONS AND FRACTALS. - ISSN 0960-0779. - STAMPA. - 192:(2025), pp. 1-14. [10.1016/j.chaos.2024.115935]

Availability:

This version is available at: 11583/2996477 since: 2025-01-10T09:01:29Z

Publisher:

Elsevier

Published


DOI:10.1016/j.chaos.2024.115935

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



An evidence-accumulating drift–diffusion model of competing information spread on networks

Julien Corsin ^a, Lorenzo Zino ^b, Mengbin Ye ^a,*,¹

^a Centre for Optimisation and Decision Science, Curtin University, Perth WA, Australia

^b Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino, 10129, Italy

ARTICLE INFO

Keywords:

Consensus
Diffusion on networks
Complex networks
Information cascade
Polarisation

ABSTRACT

In this paper, we propose an agent-based model of information spread, grounded on psychological insights on the formation and spread of beliefs. In our model, we consider a network of individuals who share two opposing types of information on a specific topic (e.g., pro- vs. anti-vaccine stances), and the accumulation of evidence supporting either type of information is modelled by means of a drift–diffusion process. After formalising the model, we put forward a campaign of Monte Carlo simulations to identify population-wide behaviours emerging from agents' exposure to different sources of information, investigating the impact of the number and persistence of such sources, and the role of the network structure through which the individuals interact. We find similar emergent behaviours for all network structures considered. When there is a single type of information, the main observed emergent behaviour is consensus. When there are opposing information sources, both consensus or polarisation can result; the latter occurs if the number and persistence of the sources exceeds a threshold value identified in the simulations. Importantly, we find the emergent behaviour is mainly influenced by how long the information sources are present for, as opposed to how many sources there are.

1. Introduction

Processes of information spread within social networks have been increasingly and extensively studied over the past few decades [1–5]. With online communications now accounting for more human interactions than in-person interactions [6], and considering the ability for false or misleading information (mis/disinformation) to damage health [7], democratic [8], and economic [9] institutions alike, understanding how ideas spread online has now become an undeniably exigent research problem — one that is being addressed by the joint efforts of multiple scientific fields.

At its roots, the science of misinformation is much of a psychological one, with human cognition and social behaviours dominating the causes of false belief formation and adoption [10]. However, because the spread of information involves countless interactions within large populations, a mathematical abstraction of choice for its study lies in agent-based models (ABMs). In ABMs, individuals (or agents) are arranged according to a given structure (or network topology), and repeatedly interact with one another, updating their states according to mathematical rules [11]. Because of their scalability, ABMs are

proficient at bridging the gap between micro- and macro-scale phenomena: in the case of information and misinformation spread, they play a crucial role in shedding light onto how certain population-wide behaviours (e.g. information cascades) emerge from single individuals following cognitive and social processes (e.g. confirmation bias, sharing homogeneity) [1,12–14]. While such processes are often well-studied in experimental psychology, their consequences are typically hard to evaluate on larger populations due to experimental constraints. As of today, the literature on modelling misinformation spread using ABMs has at least two central streams.

One natural approach is to model the spread of misinformation as one of a disease — a parallel made even more relevant by the COVID-19 crisis recently giving rise to both biological and informational epidemics, termed infodemic [15,16]. In such ABMs, when an agent has multiple network neighbours, the influence of each connection is independent. Multiple interactions, and interactions repeated over time, increase the probability of information spreading in an additive fashion, though a single interaction suffices for information to spread. Because of this property, such models are called simple contagion models [17,18]. In contrast, it was also shown that not all spreading

* Corresponding author.

E-mail addresses: julien.corsin@curtin.edu.au (J. Corsin), lorenzo.zino@polito.it (L. Zino), mengbin.ye@curtin.edu.au (M. Ye).

¹ The work of M. Ye is funded by the Western Australian Government through the Premier's Science Fellowship Program, Western Australia, and by the Australian Government through the Office of National Intelligence, Australia (NI240100203).

processes behave according to the simple contagion paradigm [19,20]. In particular, some processes of social diffusion require an individual to be simultaneously exposed to different sources in order for said individual to adopt the contagion, which has been successfully corroborated by empirical evidence [21,22]. Since then, so-called complex contagion models have blossomed [23], where an agent must be influenced by *multiple* neighbours simultaneously in order for the spreading process to occur, and where the impact of each peer is thus nonlinear.

The approaches described in the previous paragraph, however, typically assume that the decision-making processes are affected by the current interactions established by the individuals, while the influence of past interactions is typically overlooked. So far, this notion has been explored by few works in related modelling domains, such as decision-making and geopolitics, where we find instances of past-influenced diffusion models such as the social drift–diffusion model (DDM) [24,25] and sandpile models [26]. For this reason, we propose to introduce a new model of information spread that accounts for information accumulation, and to inform its design using insights from cognitive and social psychology, specifically pertaining to memory and false belief formation.

As pointed out by psychological research, belief formation relies on both past cognition and perceived social consensus [10,27]. More precisely, it has been determined that upon being presented with a novel piece of information, individuals tend to craft their own belief by assessing that information’s compatibility with previously encountered evidence, as well as beliefs of other individuals across their social circle. Possibly because of this interplay, it has been shown that repeated exposure to a given piece of information (even from a single source) suffices to increase the probability of trusting it [28,29]. Moreover, it has been observed that false beliefs are able to interfere with an individual’s decision-making process for relatively long durations after initial exposure, which has come to be known as the so-called continued influence effect (CIE). Initially formalised in cognitive psychology [30], this phenomenon has since been extensively studied, with a focus on possible causes as well as mitigation strategies — both qualitative [31, 32] and quantitative [33].

In order to capture these psychological insights, we endow agents with memory, so that they may hold and accumulate evidence towards belief in one piece of information or another. Then, we place agents on a connected graph, and give them the ability to share this information, based on a combination of both their current state as well as that of neighbouring agents. Finally, in order to model mental corrections and memory alterations, we introduce evidence noise and decay within each agent’s internal dynamics. All of these aspects of our model are formally described in Section 2.1.

Having introduced these new agent dynamics, we run numerical simulations aiming to model simple scenarios of information spread. Namely, the introduction of single, and later competing, sources of information into an initially at-rest network. Concretely, we model these sources as a subset of agents (which we call committed agents) constantly sharing the same type of information for a set duration. In order to study the impact of committed agents, we also define metrics computed from agent states, which help to outline various possible emergent behaviours such as network consensus, and polarisation. Finally, we test the robustness of our conclusions by varying the underlying network topology used to arrange agents and determine their communication capabilities. To this end, we use well-established network science models with varied degree distributions [34] such as Random-Regular, Watts–Strogatz [35] and Barabási–Albert [36] models, as well as various combinations thereof in the form of Stochastic Block Models.

For all of these network topologies, we find that our model produces the same set of emergent behaviours. In the absence of committed agents, or when a small number of committed agents are active for a short duration, consensus is the baseline emergent behaviour. Past a certain threshold in both the duration of commitments and number

of committed agents, we find that the emerging behaviour depends on how many information types are concurrently introduced in the network. If committed agents of only one type are introduced, consensus is still observed, and in fact occurs faster and more frequently than in the baseline case. However, if two types of committed agents are present and contributing opposite types of information, polarisation emerges. Concluding our findings, we note that, between the number of committed agents and the duration of their commitments, the latter quantity has a stronger influence on which emergent behaviour dominates model simulations. This last result in particular complements the current literature on misinformation spread modelling, wherein quantities analogous to the number of committed agents are often found to more heavily influence the emergent behaviour.

This article is structured as follows: In Section 2, we introduce the evidence-accumulating drift–diffusion model of information spread and motivate its design. Then, in Section 3, we use exemplar simulations to identify typical emergent behaviours of the model, and present metrics to quantify the emergence of said behaviours. In Section 4, we describe different scenarios of information spread, and discuss the simulation setup. Finally, Section 5 presents the main results of the analysis of these scenarios, and Section 6 summarises the main findings and introduces directions for future work.

2. Model

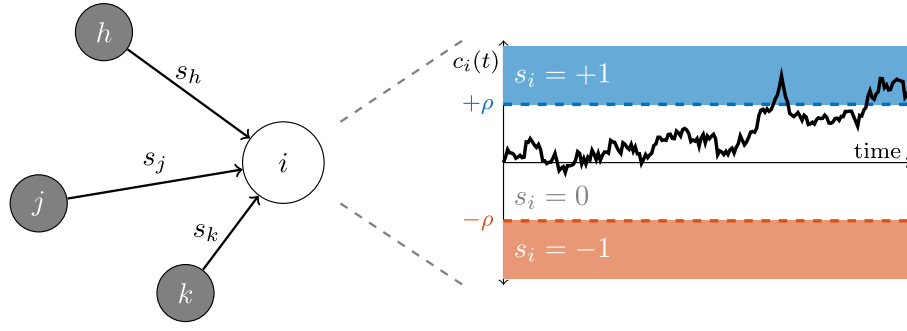
In this section, we formally introduce our evidence-accumulating model for information spread. We start by presenting the population environment and how we characterise the spread of information by means of two coupled variables per agent. Then, we introduce the dynamics that govern the evolution of these two variables. Finally, we discuss our model design and its relation with the existing literature.

2.1. Population environment

Consider a population $\mathcal{V} = \{1, \dots, N\}$ of N individuals (agents) who interact on an undirected, unweighted network structure $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the unordered edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents communication pathways between pairs of agents. Specifically, there is an edge $(i, j) \in \mathcal{E}$ if and only if agents i and j can communicate. We do not consider self-loops, i.e., $(i, i) \notin \mathcal{E}$. Accordingly, we define agent i ’s neighbourhood \mathcal{N}_i as the set of agents with whom i can communicate. That is, $\mathcal{N}_i = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\}$.

Each agent $i \in \mathcal{V}$ is characterised by two interdependent variables $s_i(t) \in \{-1, 0, +1\}$ and $c_i(t) \in \mathbb{R}$, both evolving in discrete time $t \in \mathbb{N}^* := \{1, 2, \dots\}$. The first variable $s_i(t)$ represents the stance of the piece of information being shared by agent i at time t . Specifically, we use -1 and $+1$ to represent the orientation of the information that i shares, while 0 represents the absence of any information sharing. Our model setting considers information sharing on a single topic, where agents are able to share information of opposite stances, e.g. pro/anti-vaccine narratives, left/right political agendas. Thus, $+1$ might represent information supporting a pro-vaccine narrative, in which case -1 would then represent information supporting an anti-vaccine narrative. In other words, $s_i(t)$ determines *what* agent i shares. For convenience, we will refer to $s_i(t)$ as the *stance* of agent i , even though, as discussed in the above, it more appropriately represents the stance of the information shared, since people do not always share information that is aligned with their internal beliefs.

The second variable $c_i(t)$ quantifies the amount of evidence that agent i has accumulated in favour of a particular stance. Negative values of $c_i(t)$ denote accumulating evidence in support of information with stance -1 , whereas positive values indicate evidence in support of a $+1$ -stance. The absolute value of the variable captures the amount of evidence accumulated in favour of the stance. In other words, $c_i(t)$ determines *how inclined* agent i is to share information of either type. We will refer to this variable as the *confidence* of agent i .

Fig. 1. Schematic of the dynamics for agent i .

2.2. Dynamics

We now describe the dynamics that govern the evolution and the interplay of the two variables $s_i(t)$ and $c_i(t)$. For every agent $i \in \mathcal{V}$, we define the following confidence update rule, which captures evidence accumulation and decay over time:

$$c_i(t+1) = c_i(t) + \sum_{j \in \mathcal{N}_i} s_j(t) - \delta(c_i(t) - \beta_i) + e_i(t). \quad (1)$$

Here, $\delta \in [0, 1]$ and $\beta_i \in \mathbb{R}$ are constant parameters and $e_i(t)$ is a sequence of independent and identically distributed Gaussian random variables with mean 0 and standard deviation $\sigma > 0$, independent for each agent $i \in \mathcal{V}$.

Eq. (1) poses that, at each time step t , agent i revises their confidence $c_i(t)$ according to three distinct mechanisms. Besides the current confidence $c_i(t)$, the first contribution to agent i 's revised confidence $c_i(t+1)$ is the term $\sum_{j \in \mathcal{N}_i} s_j(t)$, which accounts for the effect of information communicated by i 's neighbours on i 's confidence. Namely, agent i 's confidence is shifted in the direction of the information received from their social contacts: receiving $s_j(t) = +1$ increases agent i 's confidence in the $+1$ stance (and thus $c_i(t)$ increases), while receiving $s_j(t) = -1$ increases agent i 's confidence in the -1 stance (so, $c_i(t)$ decreases). The second contribution to $c_i(t+1)$ is the term $\delta(c_i(t) - \beta_i)$, which captures the natural decay of agent i 's confidence in any stance in the absence of evidence. Specifically, the parameter $\delta \in [0, 1]$ controls how fast confidence drifts back to the baseline value $\beta_i \in \mathbb{R}$, which can be interpreted as agent's i personal bias towards a given stance. The third and final contribution to $c_i(t+1)$ is $e_i(t)$, which is a noise term that models the imperfect nature of the information encoding process in human cognition.

After revising their confidence, each agent re-evaluates whether they wish to share information, and if so, what kind of information. Specifically, agents obey the following stance update rule:

$$s_i(t+1) = \begin{cases} +1 & \text{if } c_i(t+1) > \rho, \\ -1 & \text{if } c_i(t+1) < -\rho, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $\rho > 0$ is a fixed parameter. In other words, agent i shares information if and only if the magnitude of their confidence exceeds a threshold of ρ . In such a case, the shared value is simply the sign of the agent's confidence, as illustrated in the schematic in Fig. 1.

Observe that the two variables characterising the state of each agent co-evolve and are deeply intertwined. In fact, an agent's confidence depends on the stances received from neighbouring agents through the term $\sum_{j \in \mathcal{N}_i} s_j(t)$, while their own stance is, ultimately, a quantisation of their own confidence. All model variables and parameters associated with an agent are summarised in Table 1.

Table 1

Variables and parameters of the model.

Symbol	Meaning
$s_i(t)$	Stance of individual i at time $t \in \mathbb{N}$
$c_i(t)$	Confidence of individual i at time $t \in \mathbb{N}$
β_i	Personal bias of individual i
δ	Confidence decay rate
σ	Standard deviation of confidence noise
ρ	Sharing threshold

2.3. Discussion on the model and related literature

The core components of our model design are that of evidence accumulation (captured by $\sum_{j \in \mathcal{N}_i} s_j(t)$), and confidence decay (captured by $\delta(c_i(t) - \beta_i)$), both of which guide agent i 's information sharing process. Since information sharing is, at its core, a decision-making problem — where the decision in question pertains to both *whether* and *what* to share, we motivate our modelling choice by invoking the psychological accounts of decision-making and belief formation.

First, we observe that the role of evidence accumulation in decision-making has previously been studied by a lineage of cognitive psychology works [37,38], as well as extensively modelled [39,40]. In particular, a tool of choice for the study of evidence accumulation in decision-making is the drift–diffusion model (or DDM) [41], where evidence takes the shape of external stimuli continuously providing individuals with data supporting one of several alternatives. This accumulation of evidence continues until the amount of evidence for one of the alternatives crosses a threshold, thus completing the decision-making process by selecting said alternative. In our model, Eq. (1) is based on one particular implementation of the DDM [24], where evidence accumulates not only via external stimuli, but also via *social influence*, a process which has been shown to shape opinion formation processes in social networks [12], leading one to conjecture that social drift–diffusion models could be applied to (mis)information spread modelling. Indeed, such a paradigm has been recently used to model the spread of information [42]. However, this work differs from ours. In our model, agents share ternary values according to a quantised drift–diffusion process. In [42], a DDM is fitted to real-world experimental data so as to derive a reproduction number for a simple contagion model, which predicts the evolution of a binary set of states (share/not share). Ultimately, it is this simple contagion model (an instance of a Susceptible–Infected–Removed epidemic model), rather than a DDM, which is applied to model a process of information spread across the entire network.

Second, we note that the pervasiveness of misinformation has also been attributed in large part to its persistence in memory, as well as resistance to corrections [27]. This in turn means that even retracted (mis)information can have a continued influence on the decision-making of individuals [30]. In fact, some researchers suspect that conflicting narratives can actually coexist in memory with different




Fig. 2. Temporal evolution of the fraction of agents sharing +1 (solid blue line) and -1 (dashed orange line) for different simulation instances. All panels used the following parameters: a random-regular network with $N = 100$ agents, each with degree 4, with $\rho = 5.0$, $\delta = 0.01$, $\sigma = 0.17$, and $\beta_i = 0.0$ for all $i \in \{1, \dots, 100\}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

strengths of representation [10] — which is precisely what the sign and magnitude of $c_i(t)$ model. Consequently, modelling the varying persistence of information representations in memory, we introduce a gradual decay of evidence at a rate δ , towards a baseline evidence level β_i . This way, varying the strength of δ is conceptually equivalent to changing how long new evidence influences the decision-making process of agents.

3. System’s emergent behaviours

We start presenting our results by illustrating the range of real-world emergent behaviours associated with the spread of information and misinformation that our model is able to reproduce. Specifically, we begin by simulating the model dynamics for a fixed set of model parameters and initial conditions; the precise values of which are not important here, as we focus on qualitative behaviours in order to understand the different emergent behaviours of the model. To this end, different simulation instances of the same set of parameters and initial conditions are shown in Fig. 2, illustrating four key emergent behaviours that are typical of this system. For completeness, we note that there is a fifth (trivial) behaviour, where all agents keep a zero stance at all times (i.e., $s_i(t) = 0$ for all $t \geq 0$). Due to its simplicity, this behaviour is not illustrated by a figure.

In Fig. 2(a), we observe an initial short period where almost no information is shared, after which almost all agents continuously adopt the same stance (+1 in this example). This transition from an all-zero to an all-one stance landscape being quite sudden and short-lived (≈ 30 time steps in this example) draws a natural parallel with real-world information cascades, characterised by the viral spread of a contagion through a network [43]. Throughout this article, we will refer to an information cascade occurring whenever there is a set of connected agents, all with equal nonzero stances, and the size of this set grows over time while the set remains connected. In the event that the cascade continues until the whole agent population shares the same stance, we say that (stance) *consensus* has been reached. However, as Fig. 2(b) demonstrates, such a consensus can sometimes be reached after a substantial delay (2000 steps in this example). For this reason,

we further nuance this by distinguishing *early consensus* (Fig. 2(a)) from *delayed consensus* (Fig. 2(b)).

In Fig. 2(c), a consensus is also asymptotically reached, but differently from Fig. 2(a) and Fig. 2(b) in two aspects. First, after agents begin to share (just before $t = 2000$), we see that a nontrivial number of agents have a stance of +1 in addition to the majority of agents who have stance -1. Second, achieving consensus (to -1 in this example) is even slower (after $t > 3000$ in this example) than the delayed consensus featured by Fig. 2(b). Thinking again in terms of information cascades, one could argue that what is observed here is the presence of two candidate cascades of opposite types (-1 vs. +1) competing for network dominion, with one of them eventually winning, leading to stance consensus. We term this behaviour *perturbed consensus*, distinguishing it from *delayed consensus* because of its transient dynamics, which feature sizeable proportions of agents with both kinds of nonzero stances. However, information cascades of opposite types do not necessarily result in one type of information going extinct: in Fig. 2(d), there is again a sudden shift from no information sharing to agents sharing both kinds of information (at $t = 1000$ in this example), but this situation remains stable over the rest of the simulation time window. Because it is reminiscent of situations of social diffusion where two alternatives coexist in a network, we shall refer to this emergent behaviour as *sustained polarisation*.

3.1. Metrics to classify the emergent behaviours

The differences between the four non-trivial emergent behaviours identified above can easily be observed in a qualitative manner, as was done by plotting the stance profiles of various simulations in Fig. 2. However, given the stochastic nature of the dynamics, different behaviours can emerge from using the same model parameters (as exemplified by all plots from Fig. 2 being generated with the same set of model parameters). To account for this when quantifying the emergence of given behaviours, the main results of this paper are obtained via a large number of Monte Carlo simulations. Because of this, we develop a systematic and quantitative approach to differentiate the aforementioned emergent behaviours, motivating us to now introduce

a set of population-level metrics.

In what follows, consider a population of $N \in \mathbb{N}$ agents, whose states obey the dynamics governed by Eqs. (1) and (2) for a fixed number of $T \in \mathbb{N}$ time steps.

For each agent $i \in \mathcal{V}$, define a three-dimensional vector $x^i = [x_1^i, x_2^i, x_3^i]$, with $x_1^i := \frac{1}{T} |\{t \leq T : s_i(t) = +1\}|$, $x_2^i := \frac{1}{T} |\{t \leq T : s_i(t) = -1\}|$, and $x_3^i := \frac{1}{T} |\{t \leq T : s_i(t) = 0\}|$. In other words, this vector measures the respective *fractions of time for which agent i shared pieces of information of either type, or did not share*. Building on this, we define the *time-wise entropy* of agent $i \in \mathcal{V}$ as:

$$X(i) := - \sum_{j=1}^3 x_j^i \log_3(x_j^i), \quad (3)$$

with the convention that $x \log_3(x) = 0$ for $x = 0$.

Next, for each time $t \leq T$, we define a three-dimensional vector $y^t = [y_1^t, y_2^t, y_3^t]$, with $y_1^t := \frac{1}{N} |\{i \in \mathcal{V} : s_i(t) = +1\}|$, $y_2^t := \frac{1}{N} |\{i \in \mathcal{V} : s_i(t) = -1\}|$, and $y_3^t := \frac{1}{N} |\{i \in \mathcal{V} : s_i(t) = 0\}|$. In other words, this vector measures the respective *fractions of agents who at time t shared pieces of information of either type, or did not share*. Again, we introduce the entropy of y^t , or *agent-wise entropy* at time $t \leq T$:

$$Y(t) := - \sum_{j=1}^3 y_j^t \log_3(y_j^t). \quad (4)$$

Because x^i and y^t are stochastic vectors, $X(i)$ and $Y(t)$ are bounded in the interval $[0, 1]$. They are equal to 1 if stances are uniformly distributed across their respective domains, and to 0 if all stances are equal over said domains. Therefore, a large value of time-wise entropy $X(i)$ characterises an agent i who frequently changes the stance of the information they share over time, while a large value of agent-wise entropy $Y(t)$ indicates a specific time t for which the stances of the information being shared are quite diverse across agents.

Furthermore, notice that $X(i)$ and $Y(t)$ are respective functions of a single agent and a single time step. However, to capture emergent behaviour at the population level, it is useful to aggregate these quantities over agents, and time. For this reason, we define the respective *temporal entropy* and *spatial entropy* of stances as

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N X(i) \quad \bar{Y} = \frac{1}{T} \sum_{t=1}^T Y(t). \quad (5)$$

In the following, we use \bar{X} and \bar{Y} to characterise each of the aforementioned model behaviours, as summarised by Fig. 3.

First, consider small values of spatial entropy \bar{Y} : As previously mentioned, this happens when stances do not substantially vary across agents. Aside from the limit scenario where $\bar{X} = \bar{Y} = 0$, this corresponds to stance *consensus* occurring, with \bar{X} quantifying how *delayed*, and \bar{Y} how *perturbed*. Second, consider large values of spatial entropy \bar{Y} : As previously mentioned, this happens when the agent population is fractured into large groups of differing stances. Should these be stable over time (i.e. should \bar{X} be small enough), then *polarisation* can be said to occur, as agents effectively adopt heterogeneous long-term stances. If not, then both entropies are large, which means that agents stances vary a lot across both agents and time. In this case, there is no easily discernible emergent behaviour at the population level, although this rarely happens for the model parameters studied in this work.

While conveniently summarising agent stances with one pair of values, \bar{X} and \bar{Y} present a few limitations. First, they quantify stance differences in such a way that the difference between not sharing and sharing is considered equally significant as the one between sharing opposite types of information. Additionally, they do not capture the network topological characteristics of these stance differences, meaning that a situation involving many non-adjacent agents with the same nonzero stances would feature high \bar{Y} and small \bar{X} , which would be interpreted as polarisation if inspected only under the lens of entropy alone.

For these reasons, we introduce a third metric capturing disagreement within the network, which accounts for topological aspects of the model dynamics. Specifically, inspired from the potential function of coordination games on networks from [44], we define *network disagreement* at time t as follows:

$$D(t) := \frac{1}{4|\mathcal{E}|} \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{N}_i} (1 - s_i(t)s_j(t)) = \frac{1}{2|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} (1 - s_i(t)s_j(t)). \quad (6)$$

In Eq. (6), the summand $1 - s_i(t)s_j(t)$ measures a notion of *pairwise disagreement* between agents i and j . Mathematically, the only values pairwise disagreement can take are 0, 1, and 2, depending on the stances of the agents considered. Specifically, pairwise disagreement is maximal whenever $s_i(t)$ and $s_j(t)$ are nonzero and opposite, minimal when they are nonzero and equal, and takes an intermediate value when either or both are zero. Because of this, agent pairs $(i, j) \in \mathcal{E}$ with $s_i(t) = 0$ and $|s_j(t)| = 1$ are considered less disagreeing than pairs with $s_i(t) = +1$ and $s_j(t) = -1$ or vice-versa (a nuance which was not captured by the entropy metrics). Moreover, because pairwise disagreement takes discrete values between 0 and 2 for any agent pair and any time step $t \in \mathbb{N}$, the network disagreement $D(t)$ itself evolves within the interval $[0, 1]$ as agents dynamics from Section 2.2 play out. In particular, note that all-zero stances, *consensus*, and *polarisation* can once again be captured, respectively so by the sequence of network disagreement converging to 0.5, 0, or a different fixed constant $D^* \in [0, 1]$. In the latter case, this value D^* is then proportional to the number of agents on the boundary of polarised agent clusters. This last property in particular helps differentiate situations of polarisation where same-stance agents are spatially separated by agents of differing stances. Finally, for the sake of convenience, we define a simulation-wide summary of network disagreement, by taking the time-average of $D(t)$:

$$\bar{D} := \frac{1}{T} \sum_{t=0}^T D(t), \quad (7)$$

which completes our metrics definitions.

4. Research questions and simulation setup

With the model defined in Section 2 and its emergent behaviours discussed and classified in Section 3, we are now ready to present our research questions and describe the methodology used to implement and simulate our model. First, we design two different scenarios of information spread for which we will investigate the system's emergent behaviours. To distinguish these scenarios, we introduce new simulation-specific parameters (meta-parameters in what follows). Second, we discuss the concrete implementation of our simulations and we provide details on the region of the parameter space explored. Third, we formulate our key research questions.

4.1. Scenarios of information spread

In order to understand how our evidence accumulation process shapes the spread of information across a network, we design and run a set of Monte Carlo simulations on various synthetic network topologies. Each set of simulations is designed so as to represent different scenarios of information spread, each differing by (i) the varying presence and (ii) the persistence of *information sources*. To quantify the former, we introduce the rational numbers $(\zeta^+, \zeta^-) \in [0, 1]^2$ representing the fractions of information sources associated with respectively the +1 and -1 stances with respect to the total population size N . To quantify the latter, we introduce the non-negative integer $\tau < T$, which represents how long these sources remain active for. In what follows, we will refer to the simulation parameters (τ, ζ^+, ζ^-) as meta-parameters, allowing us to define each of the information spread scenarios of interest as an exploration of one particular subspace of this three-dimensional meta-parameter space.




Fig. 3. Characterisation of possible emergent behaviours of the model as a function of temporal (\bar{X}) and spatial (\bar{Y}) entropy.

4.1.1. Scenario I: Single information source

In our first scenario, we model a situation where information of a particular stance is introduced into a community that is otherwise initially inactive; the information being introduced by a small number of so-called *committed agents*. This allows us to study how information may propagate when introduced for the first time into a network. For this scenario, and since the model dynamics are symmetric around $s_i(t) = 0$, we choose — without loss of generality — to model committed agents who only share information with stance +1. To this end, we set $\zeta^+ > 0$ and $\zeta^- = 0$, and compute the quantity $N^+ := N\zeta^+$. If N^+ is not an integer, N^+ is rounded to the nearest integer (with ties attributed to the nearest *even* integer). Then, we sample N^+ agents from \mathcal{V} uniformly at random; that is, each agent $i \in \mathcal{V}$ has a probability of ζ^+ to be selected as a committed agent (with an adjustment if N^+ was rounded). During the simulations, committed agents are forced to share +1 at every time step $t \leq \tau$, regardless of their confidence $c_i(t)$. Once $t > \tau$, the stances of committed agents resume updating according to the rules described in Section 2.2. For brevity, we may later refer to this scenario as the single-source case.

4.1.2. Scenario II: Opposing information sources

Next, we turn to a more complex scenario, where two groups of individuals with opposite stances compete for the public information landscape. For this, we follow a similar simulation protocol as was discussed in Section 4.1.1, except for the fact that, besides the group of N^+ committed agents who share information with stance +1, there is also a second group of $N^- := N\zeta^-$ committed agents who share information with stance -1.

Formally, we implement this scenario by setting $\zeta^+ > 0$ and $\zeta^- > 0$, and computing the associated quantities N^+ and N^- , possibly rounding them as per the procedure outlined in Section 4.1.1. Then, we sample $N^- + N^+$ agents from \mathcal{V} uniformly at random; that is, each agent $i \in \mathcal{V}$ has a probability $\zeta^+ + \zeta^-$ to be selected as a committed agent. During the simulations, N^- of these newly selected agents are forced to share -1, while N^+ of them are forced to share +1 at each time step $t \leq \tau$. Thus, the probability of an agent being committed with stance +1 (resp.

-1) is ζ^+ (resp. ζ^-), with a small adjustment if there was rounding. As above, for $t > \tau$, all committed agents revert to the usual dynamics. For brevity, we may later refer to this scenario as the dual-source case.

4.2. Simulation details

We use Julia [45] to implement Monte Carlo simulations of the aforementioned scenarios. The full simulation source code is provided at [46] and a high-level description of the corresponding algorithm is provided hereafter.

First, the stances of non-committed agents $i \in \mathcal{V}$ are initialised to $s_i(1) = 0$, while those of committed agents are forced to $s_i(1) = \pm 1$, according to the scenario of information spread considered (see Section 4.1). The confidences of all agents $i \in \mathcal{V}$ are initialised to $c_i(1) = \beta_i$. Following this, all agents update their state variables exactly T times using Eq. (1) followed by Eq. (2), except for committed agents which do not update their stance until τ steps have passed. At each time step $1 \leq t \leq T$ of the simulation process, we gather the simulated agent stances and confidences, which we use to compute the metrics defined in Section 3.1. Finally, in order to obtain valid Monte Carlo approximations of said metrics, each simulation is repeated a total of R times, in order to obtain a sample mean and standard deviation of each metric — which are the final metric values presented throughout the results in Section 5. Note that committed agents are sampled at the beginning of each simulation, independently of other simulations. Since the set of preliminary simulations presented in Fig. 2 suggests that few as $T = 3000$ time steps are sufficient for any of the emergent behaviours of our model to stabilise, we elect to stop simulations after reaching this many iterations. Additionally, we fix the number of Monte Carlo samples to $R = 1024$, which we found to be sufficient for the Monte Carlo sample standard deviations of all metrics (\bar{X} , \bar{Y} , \bar{D}) to be negligible.

4.2.1. Parameter space

In the Appendix, we derive some analytical insights into the model, allowing us to restrict the parameter space to a sub-space of interest for

Table 2

Reduced parameter ranges. Bold definitions denote free parameters. Italic definitions denote fixed parameters. Definitions in roman font denote dependent parameters. The last two parameters are associated with the network structure and are changed in the last set of simulations only.

	Values	Definition
τ	$\mathbb{N}_{\leq T}$	Duration of agent commitment (steps)
ζ^+	$[0, 1]$	Fraction of committed agents with stance +1
ζ^-	$[0, \zeta^+]$	Fraction of committed agents with stance -1
ρ	$K + 1$	Sharing threshold
δ	0.01	<i>Confidence decay rate</i>
β_i	0	<i>Personal bias of individual i</i>
σ	0.2	<i>Standard deviation of confidence noise</i>
N	{100, 1000}	<i>Number of agents in the network</i>
K	4	<i>Average degree of the network</i>

our studies. First, we choose to remove any bias by setting $\beta_i = 0$ for all agents $i \in \mathcal{V}$. In fact, from our analytical observations summarised in Eq. (A.2), varying $\beta_i \neq 0$ would likely have a similar effect as varying ζ^+ and ζ^- on long-term model dynamics.

Second, since the scope of our simulations is to explore the impact of the meta-parameters τ and ζ^+ , we fix the remaining model parameters to values which do not substantially determine the overall emergent behaviours observed. In our analytical computations from Eq. (A.5), we observe that to avoid pathological patterns in how individuals share information (e.g. sudden stance transitions from -1 to +1), we need to enforce the threshold ρ to be larger than agents' degrees. For this reason, if we define K to be the average degree of the network, we fix $\rho = K + 1$. Moreover, since we do not want confidence noise to be the main contributor to the formation of information cascades, we fix $\sigma = 0.2$ based on Eq. (A.5). This way, the impact of a change to neighbouring stances on an agent's confidence is reasonably unlikely to be eclipsed by noise alone. Finally, we fix $\delta = 0.01$ so that the equilibrium point of agent confidences in Eq. (A.2) remains far enough from zero, thus enabling agent stances to remain in a (non-zero) steady state with high probability once consensus is established. All of these adjustments to the parameter space are summarised in Table 2.

4.2.2. Network topologies

Our agents interact on a network structure, which determines the set of neighbours from which an agent can receive information (as per Eq. (1)). For the underlying network topology used in our simulations, we consider four different types of network topologies, briefly described hereafter.

Random Regular (RR) networks. In RR networks, the set of agent connections \mathcal{E} is sampled uniformly at random from the set of all possible configurations in which every agent has the same constant number of K neighbours. That is, $|\mathcal{N}_i| = K \in \mathbb{N}$ for every agent $i \in \mathcal{V}$.

Watts–Strogatz (WS) model. Part of a broader category of “Small-World” networks, this network topology is obtained by first generating a network where N nodes are arranged in a circular fashion, each connected to their closest $\frac{K}{2}$ neighbours on each side. Then, for each node, each connection starting from the node is rewired with probability $\eta \in [0, 1]$, independently of the others. The rewiring is such that the connection is replaced by another one, starting from the same node but connecting it to another node selected uniformly at random from the remaining $N - 2$ nodes in \mathcal{V} . In this work, we consider a rewiring probability of $\eta = 0.1$. This process results in networks with several long-range interactions that reduces the average distance between nodes, yielding the so-called small world behaviours that is typical of many real-world social communities [47].

Barabási–Albert (BA) model. Part of a broader category of “Scale-Free” networks, these constructions are generated by iteratively adding agents to a fully-connected set of agents according to a preferential attachment rule. Specifically, each new agent chooses $\frac{K}{2}$ existing agents, with probability proportional to the current degree of existing agents, yielding a rich-get-richer phenomenon. This process results in networks possessing a power-law distribution of degrees — a property that is also observed (to some extent) in many real-world social networks [48].

Stochastic Block Model (SBM). This model combines two existing networks (or sub-communities) by pairing agents from either network together with some fixed probability λ . In this work, we consider combinations of RR and BA networks, and a link probability of $\lambda = \frac{1}{20}(\frac{K}{N} + \frac{K}{N}) = 0.004$, which is negligible with respect to both of these networks' respective average degrees.

4.3. Research questions

We start our exploration of the model by fixing the network topology, and studying what behaviours emerge for each scenario of information spread described in Section 4.1, when the model parameters are varied within the ranges of Table 2. To perform such a study, we use Random-Regular (RR) networks with $N = 100$ agents and $K = 4$. The choice of using RR networks is motivated by two main reasons. First, in RR networks all agents have the same number of neighbours, reducing the potential confounding of the population heterogeneity on the characterisation of the emergent behaviours. Second, the randomness in the network formation process allows us to introduce some randomness in the placement of information sources, providing robustness to our findings.

Once baseline results have been established for RR networks, we repeat all our simulations on more complex network topologies in order to examine the role of the network and the robustness of our findings. First, we start by investigating the scalability of the dual-source scenario by re-running the simulations from Section 4.1.2 on larger ($N = 1000$) RR networks. Next, we repeat the dual-source scenario on the other synthetic network topologies described above.

5. Results

In this section, we explore the research questions described in Section 4.3 by means of Monte Carlo simulations, as described in Section 4.2. We proceed in the same order as when introducing the different scenarios of information spread. Namely, we first investigate the single-source case, followed by the dual-source case, before studying how our observations for the dual-source case generalise to larger and different network constructions.

5.1. Single information source

First, we examine the effect of a committed minority of agents sharing information of the same stance on a RR network, following the simulation setup described in Section 4.1.1. To this end, we turn to Fig. 4, which shows measurements of $(\bar{X}, \bar{Y}, \bar{D})$ for a region of interest of the single-source meta-parameter space, namely the (τ, ζ^+) plane. For \bar{Y} , note that the colour scale used in Fig. 4(b) differs from the other panels, so as to display negligible values in a visible way.

From the heatmaps of temporal entropy \bar{X} (Fig. 4(a)) and average network disagreement \bar{D} (Fig. 4(c)), we identify two observations (or sub-scenarios) of interest, conditional on joint values of τ and ζ^+ . First, we discuss the case of *weak agent commitments*, that is, commitments which are short-lived (low τ) and involve few committed agents (low




Fig. 4. Results of the Monte Carlo simulations for a RR network in the single-source case ($\zeta^- = 0$). In panels (a), (b), and (c), we report the values of the temporal entropy \bar{X} , spatial entropy \bar{Y} , and average network disagreement \bar{D} , respectively, using a colour code, for different agent commitment durations τ and fractions of committed agents ζ^+ . In panels (d–f), we represent the temporal evolution of the network disagreement $D(t)$ for three pairs of meta-parameters (τ, ζ^+) of interest, represented by the boxed areas in panel (c).

ζ^+). Then, we discuss the case of *strong agent commitments*, i.e. commitments which are either long-lived (high τ) or involve a large fraction of committed agents (high ζ^+). Finally, we conclude our analysis of the single-source scenario by comparing the independent influences of the duration of commitment τ and the fraction of committed agents ζ^+ to determine which of the two has a stronger impact on the emergent behaviours observed for this scenario.

Weak agent commitments. When committed agents are few, and their commitments short-lived, most simulations yield delayed consensus, though we also observe a few instances of perturbed consensus. The first of these observations can be inferred by looking at the region of Fig. 4(c) and Fig. 4(a) with $\tau < 5$, which respectively feature a low average network disagreement ($\bar{D} < 0.2$) and a high temporal entropy ($\bar{X} \gg \bar{Y} \approx 0$); both of which we have previously identified in Section 3.1 as markers of delayed consensus. The presence of rare instances of perturbed consensus, on the other hand, can be inferred from the negligible but nonzero values of spatial entropy in Fig. 4(b) for $\tau < 5$, as well as the relatively slow ($t > 1000$) convergence of $D(t)$ to 0 on Fig. 4(d).

We conjecture that the delayed consensus observed for weak agent commitments is a consequence of information cascades having a very

low probability of forming within the simulation time window. In order to understand why this would be true, consider that in our model, information cascades form if a sufficiently large set of connected agents concurrently maintain their absolute confidences above the sharing threshold ρ (and thus are sharing), for a sufficiently long time to start a feedback loop in Eq. (1). Since they both facilitate the occurrence of these conditions, evidence noise and committed agents are both possible motors of cascade formation, and thus consensus. However, observe that delayed consensus is the dominant emergent behaviour in the absence of any committed agents ($\tau = 0$ or $\zeta^+ = 0$) as well as for weak agent commitments. This suggests that the main contributor to an information cascade within the simulation time window is unlikely to be the presence of committed agents, but rather evidence noise (notice that with $\sigma > 0$, a cascade forms almost surely as $t \rightarrow \infty$). However, recall that the model parameters were chosen in Section 4.2.1 so as to minimise the impact of evidence noise on cascade formation. Consequently, within our simulations, noise-induced information cascades occur with a low probability; expected time to cascade formation, and in turn the average time to consensus, are therefore large.

Then, to explain why perturbed consensus is sometimes observed, recall that the distribution of evidence noise is centred around 0. In particular, this means that information cascades associated with

both -1 and $+1$ stances have a low probability of being induced by evidence noise. In turn, it is unlikely for two information cascades to form simultaneously, making it likely that there is a significant time gap between (noise-induced) information cascades of opposite stances. Because of this, and because of the regularity of RR networks, if both $+1$ and -1 cascades form during a simulation, one of them will involve significantly more agents than the other by the time both information cascades meet within the network. However, at the interface of these information cascades, agents are more likely to convert their stance to that of the largest cascade as time progresses (a mathematical intuition for this fact is available under Eq. (A.2)). Because of the size difference between the two cascades, this conversion process is unlikely to halt until all agents are eventually of the same stance, thus explaining the observed instances of perturbed consensus.

Strong agent commitments. Increasing either of the commitments duration, or committed agents fraction, seems to have a two-fold effect on the aforementioned behaviours. First, the time needed to reach consensus is drastically reduced. Second, the previously rare instances of perturbed consensus are now nonexistent. The former effect can be seen in the fast decreasing temporal entropy \bar{X} and average network disagreement \bar{D} of Fig. 4(a) and Fig. 4(c) respectively, as the commitment duration increases above $\tau = 5$. The latter effect is characterised by the now-zero spatial entropy \bar{Y} for $\tau > 5$ in Fig. 4(b).

Analogously to the weak agent commitments case, we conjecture that these observations come as a consequence of strong agent commitments yielding a much higher probability of forming a $(+1)$ information cascade within the simulation time window — which in turn yields shorter average times to consensus. Building upon the explanation given for weak agent commitments, we note two straightforward ways to increase the chance of a $+1$ cascade forming. Specifically, one can either reduce the size of the cluster of agents that are consistently sharing, needed to initiate a positive confidence feedback loop in Eq. (1), or reduce how long this cluster needs to share for. Effectively, strong agent commitments achieve either, for the former is achieved by adding more committed agents, while the latter is achieved by increasing commitment length.

Finally, since in this scenario, only committed agents with stance $+1$ are introduced, the probability of information cascades of stance -1 forming does not increase when varying either meta-parameter. Consequently, in this sub-scenario, noise-induced -1 cascades still form with low probability. This in turn means that such cascades are unlikely to form early enough to grow to an extent which would enable them to compete with early-formed $+1$ cascades. Because of this, perturbed consensus is no longer observed in the simulations.

Respective influences of τ and ζ^+ . As noted before, for the scenario at hand, increasing either of the duration of commitments or the number of committed agents yields faster and unperturbed consensus. Consequently, one question of interest would pertain to which of these two influences is the strongest. To answer this, we note that the average network disagreement \bar{D} follows a sharp monotonic transition between delayed and early consensus as the commitment duration τ increases, with a clear threshold around $\tau = 5$ for $\zeta^+ = 0.02$. Furthermore, we note that this threshold decreases as the fractions of committed agents ζ^+ increase, as shown by the receding boundary between non-zero and zero disagreement values in Fig. 4(c), progressively going from $\tau = 5$ to $\tau = 1$ as ζ^+ increases.

However, this overall disagreement decrease seems to be much sharper in the direction of increasing τ than it is in the direction of increasing ζ^+ , suggesting that the duration of commitments plays a more significant role than the number of committed agents in facilitating an information cascade that leads to consensus. In fact, we surmise that this critical role of τ comes as a consequence of the evidence accumulation and decay mechanisms, which are typically absent in existing models of contagion. In particular, for complex contagion models [49] and independent cascade models [29], the increased spread

of contagion is linked to increasing the number of committed agents but not necessarily to increasing the length of commitments. In simple contagion models [50], the key role is played by the basic reproduction number (parameters associated with the likelihood of contagion transmission and recovery), but contagion spread is not substantially influenced by neither the number nor persistence of infection sources.

As a summary of this first simulated scenario, we conclude that the introduction of a single type of information amidst the agent population overwhelmingly yields consensus. Both longer commitments and (marginally) increasing the amount of committed agents favour the fast emergence of consensus, and decrease the number of instances of perturbed consensus.

5.2. Opposing information sources

Next, as described in Section 4.1.2, we study the effect of information sources of opposite types being concurrently introduced in a RR network. Following the single-source case, a natural question to ask pertains to how the introduction of a new type of agent commitment affects the emergent behaviours previously observed in Section 5.1. To answer this, and specifically to facilitate a comparison between the single- and dual-source cases, we find it useful to start our analysis by focusing on a sub-scenario featuring equal amounts of both types of committed agents ($\zeta^+ = \zeta^- := \zeta^\pm$). Thanks to this dimensional reduction, Fig. 5 is able to display the same metrics as Fig. 4, for the same meta-parameter region. From this, we see that Fig. 4 and Fig. 5 present similar regions of interest in terms of the meta-parameters considered, allowing us to repeat our split analysis of the results into cases of weak and strong agent commitments.

Weak agent commitments. For low values of τ and ζ^\pm , all of our measurements bear a strong similarity to the single-source scenario. In particular, Fig. 4(d) and Fig. 5(d) essentially present the same evolution of $D(t)$ for the same region of the (ζ^+, τ) plane, and the heatmaps of $(\bar{X}, \bar{Y}, \bar{D})$ do not differ significantly in their low τ and low ζ^+ regions between Fig. 4 and Fig. 5. Thus, the same conclusions can be drawn for weak agent commitments of both the single- and dual-source cases: namely, that weak agent commitments yield delayed and perturbed consensus (even when two types of information are introduced in the network).

Because of this behavioural consistency between the single- and dual-source scenarios, it is natural to make sense of this new observation by extending the rationale given for weak agent commitments in Section 5.1. Specifically, by conjecturing that both -1 and $+1$ information cascades have a low (though now equal) probability of forming at any time t . In this case, competition between information cascades of opposite types remains possible. In turn, the earliest cascade to form will likely yield consensus, leading to the observed instances of delayed (and sometimes perturbed) consensus.

Strong agent commitments. If we extend our comparison with the single-source case to the domain of strong agent commitments, significant differences start to show. First, the zero-disagreement region of Fig. 4(c) has now been replaced by a region of fixed, nonzero disagreement in Fig. 5(c). Having a closer look at the evolution of $D(t)$ for this area in Fig. 5(f) reveals that the network disagreement now converges to a fixed value $D^* \approx 0.33 \gg 0$ as t grows large. Second, in addition of now being nonzero, spatial entropy \bar{Y} now dominates temporal entropy \bar{X} , as shown by comparing the high τ or high ζ^\pm regions of Fig. 5(b) and Fig. 5(a). As pointed out in Section 3.1, all of these observations suggest that polarisation is the consistent and dominant behaviour for instances of strong agent commitments. Moreover, as commitments grow stronger, this polarisation effect is made more stable over time (as shown by $D(t)$ quickly converging to a fixed value in Fig. 5(e) and Fig. 5(f)), and involves growing numbers of disagreeing neighbours (as shown by the specific value of D^* progressively increasing from a negligible value in Fig. 5(e), to a value of 0.33 in Fig. 5(f)).




Fig. 5. Results of the Monte Carlo simulations for a RR network in the dual-source case with equal proportions of committed agents of each type ($\zeta^- = \zeta^+ = \zeta^\pm$). In panels (a), (b), and (c), we report the values of the temporal entropy \bar{X} , spatial entropy \bar{Y} , and average network disagreement \bar{D} , respectively, using a colour code, for different agent commitment durations τ and fractions of committed agents ζ^\pm . In panels (d–f), we represent the temporal evolution of the network disagreement $D(t)$ for three pairs of meta-parameters (τ, ζ^\pm) of interest, represented by the boxed areas in panel (c).

Building upon a similar reasoning as before, we claim that when commitments are strong, both types of information cascades have a high probability of forming within the simulation time window. In this case, such information cascades are more likely to form within short durations of each other compared with weak agent commitments. Because of this, and because committed agents are randomly placed in the network, such cascades are also likely to form sufficiently far from each other to allow for significant concurrent growth before confrontation occurs. Just as for the single-source case, agents at the interface of such cascades will convert to match the locally dominant stance — however, given the large and comparable sizes of cascades in this sub-scenario, this conversion process is likely to now take much longer, or stop for the rest of the simulation time window (the mathematical intuition for which is provided under Eq. (A.2)). In turn, agents on the boundaries of each cascade are more likely to maintain the same (non-zero) stance for a long time, leading to stable network polarisation with a fixed network disagreement D^* .

Weak/strong commitments boundary. The aforementioned consensus and polarisation regions are separated by a variable threshold in commitment lengths, as is visible on all heatmaps of Fig. 5. Similarly to the single-source case, this threshold is as high as $\tau = 5$ for low

amounts of committed agents, and can be observed to decrease to as low as $\tau = 1$ for high ζ^\pm . Since the dominant behaviour transitions from consensus to polarisation as τ increases, and since the average network disagreement \bar{D} *monotonically decreased* with τ in the previous scenario, one would now expect \bar{D} to *monotonically increase* with τ . However, a look at Fig. 5(c) reveals that this is not exactly the case. In fact, as the commitment length τ increases, the average disagreement \bar{D} itself seems to decrease from a low value ($\bar{D} \approx .2$) which is typical of delayed consensus, to a *local minimum* ($\bar{D} \approx 0$), before increasing to finally reach its maximum value ($\bar{D} \approx .33$) which is typical of polarisation. Interestingly, this non-monotonic transition seems to happen no matter the fraction of committed agents at play. A closer look at Fig. 5(e) suggests for the transitory behaviour to be a perturbed kind of polarisation, where the definitive boundaries of polarised groups are not decided (slowly decreasing $D(t)$) for a significant time, before fully settling on a very low number of connected disagreeing agents ($D(t) = D^*$), hinting at a quasi-consensus situation. Interestingly, this gradual decrease of $D(t)$ is much faster than the one leading to delayed consensus for weak agent commitments — meaning that no matter how many committed agents of both types are introduced, there exists an optimal commitment length which allows for a fast quasi-consensus.




Fig. 6. Results of the Monte Carlo simulations for a RR network in the dual-source case with unequal proportions of committed agents of each type ($\zeta^- \neq \zeta^+$). In each panel, we report the values of the average network disagreement \bar{D} using a colour code. In panel (a), the commitment duration is fixed to $\tau = 1$, while panel (b) focuses on the case where $\tau = 5$.




Fig. 7. Temporal evolution of the network disagreement $D(t)$ for a large ($N = 1000$) RR network computed via Monte Carlo simulations. Panel (a) represents a weak-commitment region in the single-source case, while panel (b) represents a strong-commitment region in the dual-source case with equal proportions of committed agents of both types.

Unequal proportions of information sources. For completeness, we still quantify the impact of committed populations of unequal sizes. To this end, we turn to Fig. 6, which focuses on measurements of the average network disagreement for unequal proportions of committed agents and fixed commitment durations.

In this sub-scenario, the main observation is two-fold: (i) the emerging behaviour seems to be symmetric with respect to the axis of equal proportions of committed agents of each type ($\zeta^- = \zeta^+$) (ii) straying from this axis decreases disagreement, which we interpret as facilitating consensus. In order to explain this, we briefly note that our interpretation above involving competing information cascades extends to cover this sub-scenario in a straightforward manner.

Concluding this set of simulations, we find that the introduction of two types of information sources in the agent population still allows consensus for weak agent commitments — but that stronger commitments instead introduce network polarisation. Moreover, the robustness of this polarising effect increases with longer and longer commitments, which serves as a good baseline observation to describe the quantitative differences observed between Random-Regular networks and other topologies.

5.3. Impact of the network topology

In this section, we extend the results obtained for RR networks to other network topologies, to shed light on the impact of the network structure on the emergence of the different behaviours of interest described above. Following the methodology described in Section 4.2.2, we first consider the dual-source scenario on RR networks with larger agent populations. Then, we consider different network topologies, namely the BA, WS, and SBM networks described in Section 4.2.2.

5.3.1. Larger RR networks

For the sake of brevity, and because results do not differ significantly for this sub-scenario, Fig. 7 only shows network disagreement $D(t)$ for two specific pairs of meta-parameters. Specifically, Fig. 7(a) displays the same weak-commitment region of the single-source information scenario as Fig. 4(d), while Fig. 7(b) focuses on the strong-commitment region of the dual-source information scenario shown in Fig. 5(f).

In both cases, the evolution of $D(t)$ is found to be similar as that observed for smaller RR networks, up to minor quantitative differences. For weak agents commitments (Fig. 7(a)), delayed consensus can still be inferred to occur, but the mean time needed to achieve $D(t) \rightarrow 0$ is much shorter than in smaller networks ($t \approx 300$ vs $t \approx 1700$). For strong agent commitments of opposite types (Fig. 5(f)), polarisation is also still observed, however $D(t)$ seems to converge to a higher final value ($D^* \approx 0.19$ vs $D^* \approx 0.16$).

In order to explain the former observation, note that for the same value of ζ^+ , there are ten times more committed agents for networks with $N = 1000$ than for networks with $N = 100$. Consequently, and because of our intuition on how information cascades are formed, it would make sense for the probability of an information cascade forming at any time to also follow a proportional increase in larger networks, which would explain the shorter times to consensus observed. Moving on to the latter observation, recalling that D^* measures a proportion of how many pairs of disagreeing agents are connected by an edge also makes for an intuitive explanation of the higher values of D^* observed in fully polarised larger networks.

Thus, we conclude that results obtained for both the single-source and dual-source scenarios generalise to larger networks in a straightforward fashion, with expected quantitative differences in the main metrics considered.




Fig. 8. Average network disagreement \bar{D} for RR, BA, WS, and SBM networks in the dual-source case with equal proportions of committed agents of each type ($\zeta^- = \zeta^+ = \zeta^\pm$) computed via Monte Carlo simulations.

5.3.2. Different network topologies

Finally, we reproduce our results while considering agents who interact on diverse network topologies (namely, BA and WS networks and SBMs). Specifically, we focus on the case of competing information sources with equal fractions of committed agents of each type ($\zeta^- = \zeta^+ = \zeta^\pm$) (see Section 4.1.2). The results obtained for all of the network topologies listed in Section 4.2.2 are summarised in Fig. 8, which focuses on the average network disagreement \bar{D} as the main metric being compared across network topologies. For comparison purposes, the heatmap of \bar{D} for small ($N = 100$) RR networks is reported once again in Fig. 8(a).

First, we notice that the simulation behaviours pointed out in Section 5.2 were qualitatively replicated for all networks considered, and that they follow similar transitions for similar meta-parameter values. In particular, the non-monotonicity of average network disagreement in both increasing directions of the (τ, ζ^\pm) plane is preserved, and so is the higher influence of commitment lengths (τ) over commitment numbers (ζ^\pm). Besides these similarities, and while all networks seem to agree on a baseline \bar{D} value for weak commitments, noticeable quantitative differences can be observed for strong commitments across the different network types. In particular, RR networks seem to yield higher values of average network disagreement as compared to other networks, while SBMs exhibit mixed values.

In fact, both of these differences can be readily explained by considering the defining characteristics of each network topology at hand. Because of the uniform degree distribution of RR networks, when polarisation occurs, the boundary between polarised clusters of agents is likely to always involve the same (high, perhaps maximal) number of agents, thus yielding a high D^* value. By contrast, this phenomenon is unlikely to occur in other topologies, where the degree distribution is not as uniform and information cascades are not guaranteed to split the network across a path where a lot of agents lie. In fact, considering a situation where the interface between the two polarised agent clusters comprises only one or two agents yields a rationale for the simulation behaviours observed in SBMs. Indeed, because the link probability λ between the two sub-communities of such networks is negligible, in

the extreme event of a -1 and a $+1$ cascade each forming in one of the two sub-communities, the network would likely end up polarised along the (few) links between said sub-communities, leading to an unusually low D^* value. Thus, for SBMs, the specific observed value of D^* in any simulation instance highly depends on where information cascades first form in the network.

As a conclusion, we report that the main results presented throughout Section 5.1 and Section 5.2 seem to generalise quite well to a diverse range of network topologies, despite minor quantitative discrepancies which can be explained away by differences in individual degree distributions. In particular, our key observation of the length of commitments having a higher influence than the number of committed agents on consensus/polarisation still holds, which further consolidates the hypothesis of this feature being a distinguishing consequence of introducing evidence accumulation in our model dynamics.

6. Conclusions

In this paper, we have introduced a model of competing information spread, that focused on evidence accumulation, using a social drift-diffusion model. Furthermore, we have developed information-theoretical and graph-based metrics to characterise emergent behaviours of the system, notably polarisation and consensus. Computing these metrics by means of Monte Carlo simulations, we have shown that the emergence of consensus or polarisation is conditional on the type, number, and persistence of information sources being introduced among the agent population. However, contrary to current models of information spread, we have found that out of these three influences, information source persistence is the stronger factor in deciding which behaviour emerges. We suspect that this is due to the evidence accumulation mechanism in our model, which adds a strong dependence of the dynamics on past agent interactions, absent from most current models of information spread.

As leads for future works, we propose natural generalisations of our model; namely, more realistic (but more complex) versions of the model could include heterogeneous values of δ, ρ and σ across the population,

as these quantities would realistically change on an individual basis in real-world settings. Moreover, related works considering other diffusion models have found that the location of committed agents within the network can facilitate or hinder diffusion [51–53]. We thus suggest carrying out a similar investigation for our model, by e.g. correlating the degree distribution with the distribution of committed agents, or using centrality metrics to place committed agents. Furthermore, we point out a few directions for theoretical analysis, such as to characterise the steady-state distribution of agents’ states for simple cases, such as RR networks with known locations of committed agents, and to investigate the exact (seemingly exponential) nature of the ζ^\pm/τ dependence. Finally, a rigorous real-world validation of the proposed mathematical model against real-world data should be performed.

CRedit authorship contribution statement

Julien Corsin: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Conceptualization. **Lorenzo Zino:** Writing – review & editing, Validation, Methodology, Investigation, Conceptualization. **Mengbin Ye:** Writing – review & editing, Supervision, Methodology, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Mengbin Ye reports financial support was provided by the Government of Western Australia. Mengbin Ye reports financial support was provided by the Australian Government via the Office of National Intelligence. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix. Analytical insights

To guide our exploration of the emergent behaviours of the model, we derive some mathematical insight into the dynamics, useful in restricting the model’s parameter space to a sub-space of interest for our simulations experiments, as well as in explaining how information cascades behave in our model.

Confidence equilibrium. First, consider an agent i whose neighbours $j \in \mathcal{N}_i$ have constant stances for all times $t \in \mathbb{N}$. Let N^+ denote the number of such +1 stances, and N^- denote the number of such -1 stances, so that the net stance of \mathcal{N}_i can be expressed as $\sum_{j \in \mathcal{N}_i} s_j(t) = N^+ - N^-$. Because of Eq. (1), if this net stance is nonzero, we expect the confidence of agent i to keep shifting by the value of said net stance (up to noise perturbations). Consequently, the *expected confidence* of agent i has a definite limiting behaviour, which we now derive. Rearranging Eq. (1), then taking the expected value of both sides yields the following recurrence relation:

$$\mathbb{E}[c_i(t + 1)] = (1 - \delta)\mathbb{E}[c_i(t)] + N^+ - N^- + \delta\beta_i, \tag{A.1}$$

which yields the following limit for $0 < \delta \leq 1$:

$$\lim_{t \rightarrow \infty} \mathbb{E}[c_i(t)] = (N^+ - N^-) \frac{1}{\delta} + \beta_i. \tag{A.2}$$

Thus, Eq. (A.2) implies that the long-term expected confidence of agent i is a linear function of the net stance of their neighbours. As a consequence of this, for reasonable values of model parameters (e.g. satisfying $\rho < \frac{1}{\delta} + |\beta_i|$), agent i ’s long-term stance is essentially decided by a majority rule. Conversely, this explains how a single agent can initiate information cascades in the network; namely by having a constant nonzero stance for long enough to start a confidence feedback loop between neighbouring agents. Furthermore,

Eq. (A.2) also provides insight into what can possibly happen when two information cascades of opposite stances meet in a network. As implied above, agents at the interface of such two cascades will, given enough time, convert to the stance that locally prevails in their neighbourhood. Iterating this process of successive agent conversion at the interface of the two cascades ends in two possible ways: (i) The process continues until all agents are of the same stance, yielding consensus or (ii) the process stops after a certain number of “inconvertible” neighbourhoods with zero net stance are created, yielding sustained polarisation.

Note that the duration of this conversion process highly depends on how large agent confidences have been allowed to grow prior to cascades meeting, which is intuitively analogous to how long agents have been part of said information cascades. Quantifying this further, Eq. (A.2) highlights that the maximum expected agent confidence scales with δ^{-1} . Consequently, for reasonable values of δ , any agent who has spent a lot of time partaking in an information cascade only has a negligible probability to ever revert to not sharing (or sharing opposite information) — a phenomenon which could be related to echo chambers, often observed in real-world social networks [54]

Stance continuity. The previous result implies that information cascades are formed whenever agent stances are consistent over a minimal time window and connected set of agents. However, due to the stochasticity of our model, it is theoretically possible for any agent’s stance to change solely due to evidence noise. In order to quantify how likely this phenomenon is based on model parameters, we now derive a bound on the probability that an agent’s stance takes opposite values in two consecutive time steps. To this end, we partition such an event as $A \cup B$ with:

$$\begin{cases} s_i(t + 1) = +1 \mid s_i(t) = -1 & (A) \\ s_i(t + 1) = -1 \mid s_i(t) = +1 & (B) \end{cases} \tag{A.3}$$

Starting with the first sub-event, notice that $\mathbb{P}(A)$ is maximised by agent i having an a priori confidence $c_i(t)$ of exactly $-\rho$, and all $|\mathcal{N}_i|$ of its neighbours having an a priori stance of +1. Substituting these values in Eqs. (1) and (2), and noticing that $\delta = 1$ maximises the resulting bound yields

$$\mathbb{P}(A) \leq \mathbb{P}(e_i(t) > (2 - \delta)\rho - \delta|\beta_i| - |\mathcal{N}_i|) \leq \Phi\left(\frac{|\mathcal{N}_i| + |\beta_i| - \rho}{\sigma}\right), \tag{A.4}$$

where $\Phi(x)$ is the cumulative distribution function of a standard normal distribution [55]. Applying a similar reasoning to event B , we find that the same bound also applies to $\mathbb{P}(B)$, allowing us to achieve the following bound on the total probability of our original event:

$$\mathbb{P}(s_i(t + 1)s_i(t) = -1) \leq 2\Phi\left(\frac{|\mathcal{N}_i| + |\beta_i| - \rho}{\sigma}\right). \tag{A.5}$$

Thanks to this result, the probability that a non-committed agent switches the stance of the information they share in a single time step can be made extremely small by carefully choosing model parameters. In particular, for unbiased agents ($\beta_i = 0$) in a network of average degree K , setting $\rho > K$ is enough to ensure that the average agent is unlikely to oscillate between opposite stances, in turn ensuring evidence accumulation remains central to the formation of information cascades, without being overshadowed by evidence noise.

Data availability

The repository of the code used for Monte Carlo simulations has been made public [46].

References

- [1] Vicario MD, Bessi A, Zollo F, Petroni F, Scala A, Caldarelli G, et al. The spreading of misinformation online. *Proc Natl Acad Sci* 2016;113(3):554–9. <http://dx.doi.org/10.1073/pnas.1517441113>.
- [2] Almaliki M. Online misinformation spread: A systematic literature map. In: Proceedings of the 2019 3rd international conference on information system and data mining. ACM; 2019, p. 171–8. <http://dx.doi.org/10.1145/3325917.3325938>.
- [3] Zareie A, Sakellariou R. Minimizing the spread of misinformation in online social networks: A survey. *J Netw Comput Appl* 2021;186:103094. <http://dx.doi.org/10.1016/j.jnca.2021.103094>.
- [4] Mohammadi S, Nadimi-Shahraki MH, Beheshti Z, Zamanifar K. Improved information diffusion models based on a new two-sided sign-aware matching framework in complex networks. *Chaos Solitons Fractals* 2024;187:115298. <http://dx.doi.org/10.1016/j.chaos.2024.115298>.
- [5] Zhang W, Brandes U. Conformity versus credibility: A coupled rumor-belief model. *Chaos Solitons Fractals* 2023;176:114172. <http://dx.doi.org/10.1016/j.chaos.2023.114172>.
- [6] Lieberman A, Schroeder J. Two social lives: How differences between online and offline interaction influence social outcomes. *Curr Opin Psychol* 2020;31:16–21. <http://dx.doi.org/10.1016/j.copsyc.2019.06.022>.
- [7] Chigwedere P, Seage GR, Gruskin S, Lee T-H, Essex M. Estimating the lost benefits of antiretroviral drug use in South Africa. *JAIDS J Acquir Immune Deficiency Syndr* 2008;49(4):410–5. <http://dx.doi.org/10.1097/qai.0b013e31818a6cd5>.
- [8] Bennett WL, Livingston S. The disinformation order: Disruptive communication and the decline of democratic institutions. *Eur J Commun* 2018;33(2):122–39. <http://dx.doi.org/10.1177/0267323118760317>.
- [9] Kogan S, Moskowitz TJ, Niessner M. Social media and financial news manipulation. *SSRN Electron J* 2018. <http://dx.doi.org/10.2139/ssrn.3237763>.
- [10] Ecker UKH, Lewandowsky S, Cook J, Schmid P, Fazio LK, Brashier N, et al. The psychological drivers of misinformation belief and its resistance to correction. *Nat Rev Psychol* 2022;1(1):13–29. <http://dx.doi.org/10.1038/s44159-021-00006-y>.
- [11] Macy MW, Willer R. From factors to actors: Computational sociology and agent-based modeling. *Annu Rev Sociol* 2002;28(1):143–66. <http://dx.doi.org/10.1146/annurev.soc.28.110601.141117>.
- [12] Acemoglu D, Ozdaglar A, ParandehGheibi A. Spread of (mis)information in social networks. *Games Econom Behav* 2010;70(2):194–227. <http://dx.doi.org/10.1016/j.geb.2010.01.005>.
- [13] Acemoglu D, Ozdaglar A, Siderius J. A model of online misinformation. *Rev Econ Stud* 2023. <http://dx.doi.org/10.1093/restud/rdad111>.
- [14] Davis JT, Perra N, Zhang Q, Moreno Y, Vespignani A. Phase transitions in information spreading on structured populations. *Nat Phys* 2020;16(5):590–6. <http://dx.doi.org/10.1038/s41567-020-0810-3>.
- [15] Cinelli M, Quattrociochi W, Galeazzi A, Valensise CM, Brugnoti E, Schmidt AL and others. The COVID-19 social media infodemic. *Sci Rep* 2020;10(1).
- [16] Bulai IM, Sensi M, Sottile S. A geometric analysis of the sirs compartmental model with fast information and misinformation spreading. *Chaos Solitons Fractals* 2024;185:115104. <http://dx.doi.org/10.1016/j.chaos.2024.115104>.
- [17] Nekovee M, Moreno Y, Bianconi G, Marsili M. Theory of rumour spreading in complex social networks. *Phys A* 2007;374(1):457–70. <http://dx.doi.org/10.1016/j.physa.2006.07.017>.
- [18] Tambuscio M, Ruffo G, Flammini A, Menczer F. Fact-checking effect on viral hoaxes. In: Proceedings of the 24th international conference on world wide web. ACM; 2015, p. 977–82. <http://dx.doi.org/10.1145/2740908.2742572>.
- [19] Centola D, Macy M. Complex contagions and the weakness of long ties. *Am J Sociol* 2007;113:702–34. <http://dx.doi.org/10.1086/521848>.
- [20] Watts DJ, Dodds PS. Influentials, networks, and public opinion formation. *J Consum Res* 2007;34(4):441–58. <http://dx.doi.org/10.1086/518527>.
- [21] Mønsted B, Sapiezynski P, Ferrara E, Lehmann S. Evidence of complex contagion of information in social media: An experiment using Twitter bots. *PLoS One* 2017;12(9). <http://dx.doi.org/10.1371/journal.pone.0184148>.
- [22] Sprague DA, House T. Evidence for complex contagion models of social contagion from observational data. *PLoS One* 2017;12(7). <http://dx.doi.org/10.1371/journal.pone.0180802>.
- [23] Guilbeault D, Becker J, Centola D. Complex contagions: A decade in review. In: Complex spreading phenomena in social systems. Springer International Publishing; 2018, p. 3–25. http://dx.doi.org/10.1007/978-3-319-77332-2_1.
- [24] Tump AN, Pleskac TJ, Kurvers RHJM. Wise or mad crowds? The cognitive mechanisms underlying information cascades. *Sci Adv* 2020;6(29). <http://dx.doi.org/10.1126/sciadv.abb0266>.
- [25] Fudenberg D, Newey W, Strack P, Strzalecki T. Testing the drift–diffusion model. *Proc Natl Acad Sci* 2020;117(52):33141–8. <http://dx.doi.org/10.1073/pnas.2011446117>.
- [26] Lu P, Yang H, Li M, Zhang Z. The sandpile model and empire dynamics. *Chaos Solitons Fractals* 2021;143:110615. <http://dx.doi.org/10.1016/j.chaos.2020.110615>.
- [27] Lewandowsky S, Ecker UKH, Seifert CM, Schwarz N, Cook J. Misinformation and its correction. *Psychol Sci Public Interest* 2012;13(3):106–31. <http://dx.doi.org/10.1177/1529100612451018>.
- [28] Hasher L, Goldstein D, Toppino T. Frequency and the conference of referential validity. *J Verb Learn Verb Beh* 1977;16(1):107–12. [http://dx.doi.org/10.1016/s0022-5371\(77\)80012-1](http://dx.doi.org/10.1016/s0022-5371(77)80012-1).
- [29] Budak C, Agrawal D, Abbadi AE. Limiting the spread of misinformation in social networks. In: Proceedings of the 20th international conference on world wide web — WWW '11. ACM Press; 2011, p. 665–74. <http://dx.doi.org/10.1145/1963405.1963499>.
- [30] Johnson HM, Seifert CM. Sources of the continued influence effect: When misinformation in memory affects later inferences. *J Exp Psychol [Learn Mem Cogn]* 1994;20(6):1420–36. <http://dx.doi.org/10.1037/0278-7393.20.6.1420>.
- [31] Kendeou P, Walsh EK, Smith ER, O'Brien EJ. Knowledge revision processes in refutation texts. *Discourse Process* 2014;51(5–6):374–97. <http://dx.doi.org/10.1080/0163853x.2014.913961>.
- [32] Cook J, Lewandowsky S, Ecker UKH. Neutralizing misinformation through inoculation: Exposing misleading argumentation techniques reduces their influence. *PLoS One* 2017;12(5). <http://dx.doi.org/10.1371/journal.pone.0175799>.
- [33] Ecker UKH, Lewandowsky S, Swire B, Chang D. Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bull Rev* 2011;18(3):570–8. <http://dx.doi.org/10.3758/s13423-011-0065-1>.
- [34] Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang D. Complex networks: Structure and dynamics. *Phys Rep* 2006;424(4–5):175–308. <http://dx.doi.org/10.1016/j.physrep.2005.10.009>.
- [35] Watts DJ, Strogatz SH. Collective dynamics of ‘small-world’ networks. *Nature* 1998;393(6684):440–2. <http://dx.doi.org/10.1038/30918>.
- [36] Albert R, Barabási A-L. Statistical mechanics of complex networks. *Rev Modern Phys* 2002;74(1):47–97. <http://dx.doi.org/10.1103/revmodphys.74.47>.
- [37] Lee MD, Cummins TDR. Evidence accumulation in decision making: Unifying the take the best and the rational models. *Psychonomic Bull Rev* 2004;11(2):343–52. <http://dx.doi.org/10.3758/bf03196581>.
- [38] Usher M, Tsetsos K, Yu EC, Lagnado DA. Dynamics of decision-making: from evidence accumulation to preference and belief. *Front Psychol* 2013;4. <http://dx.doi.org/10.3389/fpsyg.2013.00758>.
- [39] Usher M, McClelland JL. The time course of perceptual choice: The leaky, competing accumulator model. *Psychol Rev* 2001;108(3):550–92. <http://dx.doi.org/10.1037/0033-295x.108.3.550>.
- [40] Brown SD, Heathcote A. The simplest complete model of choice response time: Linear ballistic accumulation. *Cogn Psychol* 2008;57(3):153–78. <http://dx.doi.org/10.1016/j.cogpsych.2007.12.002>.
- [41] Ratcliff R, McKoon G. The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput* 2008;20(4):873–922. <http://dx.doi.org/10.1162/neco.2008.12.06.420>.
- [42] Alvarez-Zuzek LG, Grujic J, Gallotti R. Drift diffusion model to understand (mis)information sharing. 2024. <https://arxiv.org/abs/2401.17846>.
- [43] Jalili M, Perc M. Information cascades in complex networks. *J Complex Netw* 2017;5(1):665–93. <http://dx.doi.org/10.1093/comnet/cnx019>.
- [44] Aghbolagh HD, Ye M, Zino L, Chen Z, Cao M. Coevolutionary dynamics of actions and opinions in social networks. *IEEE Trans Autom Control* 2023;68(12):7708–23. <http://dx.doi.org/10.1109/tac.2023.3290771>.
- [45] Bezanson J, Edelman A, Karpinski S, Shah VB. Julia: A fresh approach to numerical computing. *SIAM Rev* 2017;59(1):65–98. <http://dx.doi.org/10.1137/141000671>.
- [46] Corsin J. Julia implementation of the competing information spread evidence drift–diffusion model. 2024. <https://gitlab.com/Holbormon/ciseddm>.
- [47] Newman MEJ. The structure and function of complex networks. *SIAM Rev* 2003;45(2):167–256. <http://dx.doi.org/10.1137/s003614450342480>.
- [48] Clauset A, Shalizi CR, Newman MEJ. Power-law distributions in empirical data. *SIAM Rev* 2009;51(4):661–703. <http://dx.doi.org/10.1137/070710111>.
- [49] Ye M, Zino L, Mlakar Ž, Bolderdijk JW, Risselada H, Fennis BM, et al. Collective patterns of social diffusion are shaped by individual inertia and trend-seeking. *Nature Commun* 2021;12(1). <http://dx.doi.org/10.1038/s41467-021-25953-1>.
- [50] Nowzari C, Preciado VM, Pappas GJ. Analysis and control and of epidemics: a survey of spreading processes on complex networks. *IEEE Control Syst* 2016;36(1):26–46. <http://dx.doi.org/10.1109/mcs.2015.2495000>.
- [51] Gao T, Zino L, Ye M. Effect of network structure and committed minority placement in promoting social diffusion. *IEEE Trans Comput Soc Syst* 2023;1–14. <http://dx.doi.org/10.1109/tcss.2023.3303568>.
- [52] Banerjee A, Chandrasekhar AG, Duflo E, Jackson MO. The diffusion of micro-finance. *Science* 2013;341(6144):1236498. <http://dx.doi.org/10.1126/science.1236498>.
- [53] Kempe D, Kleinberg J, Tardos E. Maximizing the spread of influence through a social network. In: Proceedings of the 9th ACM SIGKDD international conference on knowledge discovery and data mining. 2003, p. 137–46. <http://dx.doi.org/10.1145/956750.956769>.
- [54] Cinelli M, De Francisci Morales G, Galeazzi A, Quattrociochi W, Starmini M. The echo chamber effect on social media. *Proc Natl Acad Sci* 2021;118(9). <http://dx.doi.org/10.1073/pnas.2023301118>.
- [55] Ross SM, Peköz EA. A second course in probability. Cambridge University Press; 2023.