



Knowledge diffusion in social networks under targeted attack and random failure: the resilience of communities

Piergiuseppe Morone¹ · Rocco Caferra¹ · Antonio Lopolito²

Received: 14 November 2022 / Accepted: 22 June 2023 / Published online: 4 July 2023

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2023

Abstract

Knowledge diffusion is a complex and demanding process that requires coordination and collaboration between agents with different levels of knowledge, to establish fruitful learning interactions. In this paper, we develop an agent-based model to investigate how different behavioral/sociological rules can alter, strengthen, or weaken this process. We observe that, during normal times, different aggregation strategies are apparently irrelevant for determining differences in learning opportunities. However, under crisis, there is an observable outperformance of social structures with established communities, characterized by both strong ties (i.e., intense contacts within communities) and weak ties (i.e., knowledge spillover across communities). We further test system resilience, considering interruptions to the knowledge diffusion of expert agents and the random temporary removal of agents (simulating a viral outbreak). We discuss how these scenarios may explain economic phenomena and explore the implications for policies aimed at mitigating knowledge and economic inequalities.

1 Introduction

The idea that “fish travel in schools, birds migrate in flocks, and ants build trails” (Sumpter 2010) has attracted researchers from different fields, including economists. Fish travel in schools and birds migrate in flocks to better protect themselves from predators, while ants build trails to track their path to a source of nutrition. And humans? As *social animals* (Aristotle), humans are not so different from other species in the ways in which they organize themselves. Such social aggregation is fruitful for the creation and use of ideas, favoring social development (Schuller and Theisens 2010). Put differently, humans aggregate in order to evolve.

✉ Piergiuseppe Morone
piergiuseppe.morone@unitelmasapienza.it

¹ Department of Law and Economics, UnitelmaSapienza University of Rome, Rome, Italy

² Department of Economy, Management and Territory, University of Foggia, Foggia, Italy

The discovery that sociological rules may determine economic outcomes has led to a general need to rethink macrosystems, considering the coordination and interaction between parts, or human organizational patterns (Dobbin, 2004; Kirman, 2010). As a result, attempts have been made to explain economic dynamics based on different rules of agent aggregation, outlining potential paths to economic growth. In this work, a crucial role has been attributed to the diffusion of knowledge as a prime mover for economic success—or its inverse, economic inequality (Morone and Taylor 2004). From this, a question naturally arises: “Individuals may have knowledge that is useful to others, but does this knowledge flow to those who need it?” (Alatas et al. 2016).

In 2004, Cowan and Jonard claimed: “The details of who is connected to whom will clearly affect what type of information is passed, how much, and how efficiently.” From this assumption, they asked, “if the network structure is exogenous, how do the structural properties of the network affect aggregate outcomes?” (Cowan and Jonard 2004, p. 1558). In the present research, we aimed at extending the literature, focusing on social network structure and its implications for efficiency (i.e., high average knowledge), equity (i.e., low knowledge heterogeneity), and sustainability (i.e., high resilience to shock), with respect to knowledge diffusion. We kept economic features as simple as possible, considering that agents do not always follow rational and complicated optimization rules in their daily interactions, but simply acquire knowledge from their daily contacts, in accordance with different behavioral rules (as in Block et al. 2020). This style of knowledge acquisition may result in a wide range of economic outputs, under different conditions. This underscores the relevance of using agent models to study the coordination of collective phenomena and assessing sociological aspects (e.g., the characteristics of structural social capital) to understand and predict the dynamics of economic phenomena.

The remainder of the paper is organized as follows: Sect. 2 reviews the existing literature, Sect. 3 introduces the methodology and main working hypotheses, Sect. 4 reports the results, and Sect. 5 provides concluding remarks.

2 Literature review

In this section, we move from an exposition of the main results of prior analyses of knowledge networks to a discussion of behavioral rules for interaction that may shed light on the nature and dynamics of society (i.e., the ways in which interaction norms may explain the dynamics and resilience of social systems, with a specific emphasis on the human tendency to aggregate in communities). The final section of the review focuses on targeted attacks and random failures in social networks, describing realistic examples based on very recent episodes (e.g., the COVID-19 pandemic).

Similar to viral spread, knowledge diffusion occurs among interacting (i.e., communicating) subjects, and different network topologies determine different properties of knowledge invasion, spread, and persistence (Cowan and Jonard 2004; May and Lloyd 2001). However, differently from a viral diffusion scenario, in which even a single interaction can be sufficient for infection, the knowledge diffusion process

requires repeat interactions for learning to occur (Morone and Taylor 2004). Therefore, links are crucial for facilitating dissemination, and the intensity of connections might contribute to defining particular knowledge diffusion patterns.

Considering network theory, several researchers have investigated the structure and the dynamics (Liu and Zhang 2014; Greenan 2015; Snijders 2001; Snijders et al. 2010; Steglich et al. 2010) of knowledge networks, studying the spatial and temporal features of diffusion (Strang and Tuma 1993). Considering the spatial dimensions, existing studies include small-world (Cowan and Jonard 2003, 2004), scale-free (Lin and Li 2010; Tang and Mu, 2010), and local-world non-uniform hypernetworks (Yang and Hu, 2015), as well as regular networks (Zeng et al., 2019). In the present research, we aimed primarily at characterizing exogenous network formations, considering realistic social interactions. That is, we accounted for the formation of networks in which agents, following different rules of behavior, select daily contacts via intuition (as proposed in Block et al. 2020), forming dynamic networks with different properties. The idea of an exogenous network assumes that interactions are not learning-driven, but exogenous to the potential learning capabilities of specific agents (as in Cowan and Jonard 2004).

From a sociological perspective, the spread of knowledge is related to the co-existence of strong and weak ties within a social network. Strong ties refer to intense connections among two nodes, while weak ties define bridges across distant subgroups formed by stronger ties. Weak ties are particularly effective for spreading knowledge across groups, thereby reducing fragmentation within the larger network and improving diffusion (Granovetter 1973). However, strong ties more readily stimulate innovation and knowledge exchange (Nelson 1989). Previous research has studied knowledge diffusion dynamics, considering different types of intra-clique, inter-clique, and extra-clique connections (Bala and Goyal 2001; Midgley et al. 1992; Ally and Zhang 2018). When both strong and weak ties exist, a network can be characterized as complex (Kim and Park 2009). This suggests that access to diversified ties is important for improving productivity and innovation, through knowledge diffusion (Todo et al. 2016).

Strong ties can be represented by clusters and communities (i.e., subgroups of actors with close and frequent interaction). Examples of these clusters include firms involved in R&D (Rappa and Debackere, 1992), scientific communities (Lambiotte and Panzarasa 2009), and, more generally, communities of individuals who regularly exchange ideas. Such clusters may benefit knowledge diffusion and creation (Rappa and Debackere, 1992; Storck and Hill 2009). In particular, Malmberg and Power (2005) showed that clusters have been rapidly growing in both academic and policymaking circles, promoting knowledge exchange and acquisition through collaborative interaction. In this way, clusters can represent melting pots for innovation (Günther and Meissner 2017).

A proper balance of strong ties (i.e., knowledge diffusion within communities) and weak ties (i.e., spillover between communities and/or other network nodes) can guarantee access to knowledge, thereby favoring development and reducing inequalities. This is supported by social capital theories (Putnam 2001) and the notion of *structural social capital* (Nahapiet et al., 1998), which describes the benefits an individual accrues through learning from their network. Following this literature,

bonding and bridging capital (i.e., intense local interactions within a given community and cross-community interaction, respectively) might be complementary for fostering economic growth, since “the relations within and between social groups at different levels of society shape the prospects for sustainable, equitable growth” (Woolcock, 1998).

Against this background, recent events (e.g., COVID-19) have highlighted the importance of social network structure for determining collective phenomena and the relevance of social interaction to productivity (Breetze and Wild, 2022; Jialu et al. 2021). Thus, it is useful to revisit and analyze the evolution of social networks and the sociological aspects that define interaction rules. By identifying the changing network configurations of individual contact (selected via choice) and organizational routines, with regard to different types of agents (e.g., individuals, firms), we can observe how knowledge spreads under both normal conditions and conditions of stress.

Another important network feature is change in the intensity of relationships. Social ties may be strengthened, weakened, or broken due to internal and/or external circumstances. Furthermore, the network literature has underlined the importance of investigating the sustainability of knowledge networks (Zhao et al. 2020). Importantly, networks can be subjected to malicious attacks that hinder communication (Dong et al. 2013). These attacks can be random or targeted, depending on the ways in which they disable certain nodes. It is difficult to identify a network structure capable of preventing both random and targeted attacks, since the optimal configuration for defending against targeted attacks can be most vulnerable to random attacks (Zhang et al. 2016).

There are several real-life examples of targeted and random attacks. For instance, examples of targeted attacks (which affect only a subset of the network, such as experts) include high-tech firms colluding to prevent knowledge spread outside their cluster, to maintain competitiveness (Bacchiega et al. 2010); financial traders withholding privileged knowledge; and experts protecting their experiential knowledge in order to preserve a competitive advantage in a difficult job market. An example of a random attack generating behavioral change throughout the entire network (considering both experts and non-experts) is the COVID-19 pandemic. During the viral outbreak, different agents were infected and their communication with the rest of the network was temporarily interrupted. It is well documented that the interruption of face-to-face interaction, accompanied by difficulties in preserving effective online interaction, led to communication dysfunction between subjects, generating inequalities (Haelermans et al. 2022).

The literature on networks shows that learning processes within social networks are interrupted or altered by outside forces according to their degree of resilience to different attack strategies (Latif et al., 2013). While targeted attacks affect specific nodes (i.e., expert agents), interrupting communication and learning, random failures tend to disable the communication of heterogeneous agents.

Scholars have attempted to identify which network properties favor resilience and robustness (De Domenico et al. 2014; Latif et al., 2013; Li et al. 2019). However, previous analyses of network structure have failed to sufficiently consider the behavioral rules underlying sociological drivers of network dynamics. Therefore, in the present

research, we considered if—and how—certain behavioral interaction rules might foster resilience to both targeted and random attacks, in order to maintain the functioning of social networks, even in the context of adversity.

Specifically, the research aimed at contributing to the literature on social aggregation and learning patterns, with respect to their impact on growth and inequality. We considered a simulated framework with different stylized (but realistic) aggregation rules representing social interactions. Following Block et al. (2020), we modeled the strengths and weaknesses of ties based on random, community, and repetition strategies. As discussed in the “Methodology” section, in the random case, interactions were sparse and the frequency with which the same agents interacted was lower, demonstrating the predominance of weak ties. In the other cases, there was a greater frequency of strong ties: The community strategy was based on the creation of communities among agents with the highest number of shared contacts; and the repetition strategy was based on repeat interaction with the same partner over time. In the repetition case, there was a high likelihood of strong ties being formed but a lower (but non-zero) probability of weak ties being formed with nodes outside the community or set of repeat contacts.

The number of contacts was maintained as a constant across all scenarios, while the intensity of strong ties ranged from extremely unlikely (in the random case) to highly likely (in the repetition case). The community case represented an intermediary between the random and repetition cases, with frequent but diffused contact among community members. This allowed expert nodes to disseminate knowledge at an intermediate frequency (i.e., more frequently than in the random case, but less frequently than in the repetition case) to a larger number of agents.

Network structure may be considered structural social capital, defining (following Nahapiet and Ghosal, 1998) the network of interpersonal relationships shared by everyone. In our model (and beyond), knowledge was spread through social interaction between agents; accordingly, the network of interpersonal relationships (and the structural social capital within a society) defined the manner in which knowledge was spread. From this perspective, we tested the effects of an open but disorganized network (random case), an open but organized network (community case), and a very closed network (repetition case).

Societies are repeatedly affected and attacked by both external and internal conditions. In some cases, members of the society, themselves, may change behavior and thus alter the knowledge dissemination process within the network. In other cases, exogenous shocks may alter the network’s operations (as during COVID-19). In the present research, we attempted to show how the three scenarios (i.e., random, community, and repetition) lead to the same aggregate level of knowledge and knowledge distribution under conditions of internal change, but a different level of knowledge and knowledge distribution under conditions of shock.

3 Methodology

To better convey the construction of our model, the following sections clarify: (i) the ex-ante assignment of knowledge, (ii) the learning process, (iii) the analyzed interaction rules, (iv) the resilience tests, and (iv) the main performance indicators.

3.1 Knowledge distribution

It is widely known that ex-ante inequality in the endowment of knowledge contributes to determining economic performance (Morone and Taylor 2004). Considering that the distribution of knowledge can be a proxy for the distribution of wealth, we applied Pareto's law of income distribution (as observed in several countries) to determine the knowledge distribution to agents (i.e., allocating 20% of the nodes 80% of the wealth). Although the empirical evidence of this income structure is mixed (Clementi and Gallegati 2005), we applied it because it considers a starting condition with both relatively high inequality and relatively high learning opportunities. Following this, we populated the network with 1,000 agents (as in Block et al. 2020), disseminating 50,000 tokens of knowledge across the nodes. Therefore, vertexes were subdivided into two groups: 200 expert agents (i.e., 20% of the total nodes), receiving 40,000 knowledge tokens (i.e., 80% of the total tokens) in a random distribution; and 800 non-expert agents (i.e., the other nodes), receiving 10,000 knowledge tokens in a random distribution. Figure 1 reports the initial distribution of knowledge tokens. The x-axis indicates the 1000 agents sorted in descending order, from higher to lower level of knowledge (reported on the y-axis). The vertical line divides the experts (i.e., the first 200 agents) from the non-expert agents (on the right side).

3.2 Learning process

In line with previous research (Morone and Taylor 2004), we introduced the concept of *learning gains* from interactions, while simplifying the process to better evidence and isolate the effect of the different interaction rules, in accordance with our primary research aim. Cowan and Jonard (2004) developed a model capturing the

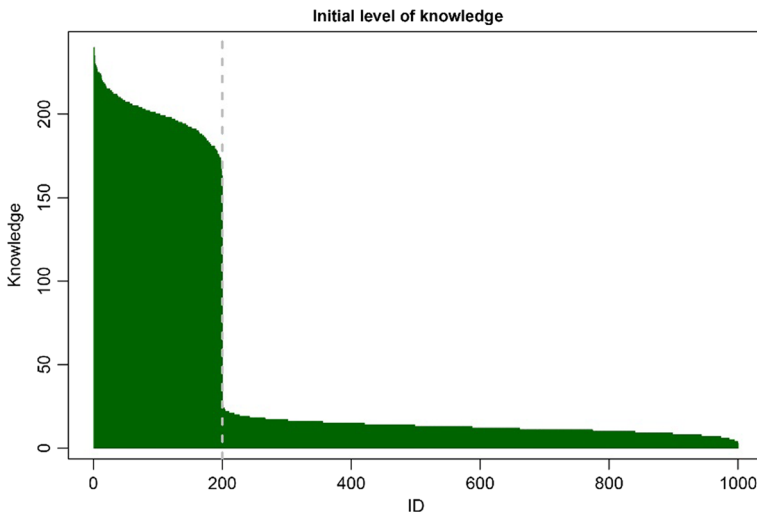


Fig. 1 Initial distribution of Knowledge

effects of incremental innovation and the diffusion of these effects throughout a network of heterogeneous agents. Even in their model, agents were endowed with different levels of initial knowledge: A small number of agents were assigned “expert” roles and attributed a higher level of knowledge. Knowledge was assumed to be exchanged through interaction between an expert (or individual with a higher level of knowledge) and an agent with less knowledge.

The main idea behind learning gains is that knowledge transfer is not obtained in a single interaction, but requires repeated contact, depending on the capacity of the receiving node. In the real world, this is evidenced by the years of study required for students to earn a degree, years of R&D required to determine innovations, and years of research required for academics to make an original contribution to the literature.

In our model, learning gains were assumed to be constant, with each vertex acquiring one token of knowledge from a more expert node, as follows: k_{it} is the level of knowledge of agent i at time t and defined as the number of tokens i owns; it changes according to the following rule, after i interacts with agent j at time $t+1$:

$$k_{it+1} = \begin{cases} k_{it} + 1 \text{ and } k_{jt+1} = k_{jt} & \text{if } k_i < k_j \\ k_{it} \text{ and } k_{jt+1} = k_{jt} + 1 & \text{if } k_i > k_j \text{ or } k_i = k_j \end{cases}$$

3.3 Interaction strategies

The spread of ideas (similar to the spread of viruses) is crucially affected by aspects of network structure, such as path length (i.e., network distance, or the number of network steps needed to connect two nodes) and level of clusterization. Different hypotheses have been proposed to describe the spread mechanisms that affect the diffusion of a new behavior throughout a network. One view is based on the *strength of weak ties* hypotheses; this perspective considers the spread of behavior as simple as that of a virus, requiring only a single interaction with an infected (or newly behaving) person (Granovetter 1973; Watts, 1999). In this case, the ideal structure to optimize rapid diffusion throughout the network is a random network with many long ties, a very short average path length, and low clusterization. This ensures that each agent is capable of reaching many others in different neighborhoods, thereby reducing redundancy at the same time.

In contrast, the *complex contagion* hypothesis (Centola and Macy, 2007) sees the diffusion of new behavior as more complex than the spread of a disease. In this hypothesis, numerous exposures to the carrier of the novelty are required for an agent to decide to adopt the new behavior. In this case, the ideal structure for diffusion is that of a highly clustered network with sufficient social reinforcement to encourage adoption.

Block et al. (2020) proposed a set of interaction strategies to represent different network conditions (according to average path length and network clusterization), with the aim of identifying the most efficient strategy to prevent viral spread.

Applying this analysis to behavioral diffusion, in the present research, we aimed at assessing the impact of different social interaction strategies on spread throughout a network.

Specifically, we analyzed three strategies: random, repetition, and community. The random strategy involved a network with weak ties, a very short average path length, and a low clusterization level, representing the least organized network structure. In this condition, agents randomly selected alters with which to interact (without following any organizational rule), within their own neighborhood. On the other end of the continuum, the repetition strategy involved a network structure based on repeated contact, where the number of partners (rather than the number of interactions) was reduced, creating so-called *social bubbles*. Within this network, knowledge was quickly disseminated between partners via the high frequency of interaction, but rarely shared outside social bubbles.

The final strategy built on the strength of communities. In this model, individuals mostly interacted with those with whom their neighbors also commonly interacted, reflecting the idea that agents tend to aggregate with those who share similar interests or occupy similar “melting pots,” creating overlapping networks of neighbors. We considered the community strategy an intermediate social aggregation pattern located between the other two extremes, which considered prevalently weak (in the random case) and strong (in the repetition case) ties. However, we did not include the *homophilistic* case (analyzed in Block et al. 2020), which could be considered a further intermediary between the random and community scenarios.

As in Block et al. (2020), we started with a network of 1,000 agents, defining a *static* neighborhood for each vertex. In this work, we only emphasize aspects that are novel to Block et al. (2020) and refer the reader to the original study for further details on the analysis and nomenclature employed.

The web of interaction opportunities among agents was determined by four subprocesses: (i) geographical proximity, (ii) group membership, (iii) homophily, and (iv) daily random meets. Considering geographical proximity, actors were randomly scattered across a 2D space, and a random number of ties (ranging from 4–12) were formed based on the Euclidean distance between points (assuming an agent density of $d_{\text{geo}}=0.3$). On average, groups had eight members, and each actor formed ties with other members at random (assuming an agent density of $g_{\text{groups}}=0.6$). The homophilistic proximity between subjects ranged from 0–100. Considering the similarity formula (i.e., closest to Euclidean distance), a random number of ties with the closest alters (ranging from 4–12) were formed. The fourth subprocess represented random z meetings/links (500 in the entire network).

The *dynamic* was provided by a multinomial choice model in which each agent i , picked at random in each simulation step, selected j from its neighborhood. The weighted probability of selecting j differed across each strategy. In the random scenario, each agent j had an equal probability of being selected. In the repetition case, j was weighted in accordance with the number of times that i met j in the three prior interactions. In the community case, the probability of selecting j was determined by a statistic indicating the number of alters shared between i and j . The higher the number of shared alters, the higher the probability that a cluster would be formed with those contacts. As in Block et al. (2020), the statistics were weighted

considering the same parameters ($s_{\text{community}}$: $\alpha=0.75$; $s_{\text{repetition}}$: $\alpha=2.5$). The final weighted probability was an exponentiated linear function of the resulting statistics for each case (see Appendix A for the computational details).

Figure 2 reports the difference in aggregation between the community and repetition cases, with the community case showing clusters with more members. As the network was large, we opted for a Fruchterman–Reingold display, with nodes visually arranged based on their network distance. This resulted in linked nodes being displayed close together.

In a similar vein, Fig. 3 illustrates the intensity of weak and strong ties across agents, considering the frequency of contact.

The random scenario resulted in the highest intensity of sparse contacts, with low-frequency contact between each pair of vertices. In contrast, the community case displayed a high level of contact with an intermediate intensity (in blue), while the repetition strategy reflected dense, high-frequency pairwise contacts, with dyads disconnected from the rest of the network.

3.4 Resilience tests

Figure 4 illustrates workflows for the different scenarios analyzed. As a reference point, we considered simple interaction among agents exchanging knowledge, without setting any kind of restriction. In this case, the simulation consisted of 500 periods, each with 1,000 micro-steps (resulting in a total of 500,000 micro-steps), with an agent I randomly selected for interaction in each step. We observed that 500 periods were sufficient to achieve a steady state. (The same ratio was applied to select the periods in the other cases.) The other two cases represented a targeted attack and random failure, respectively.

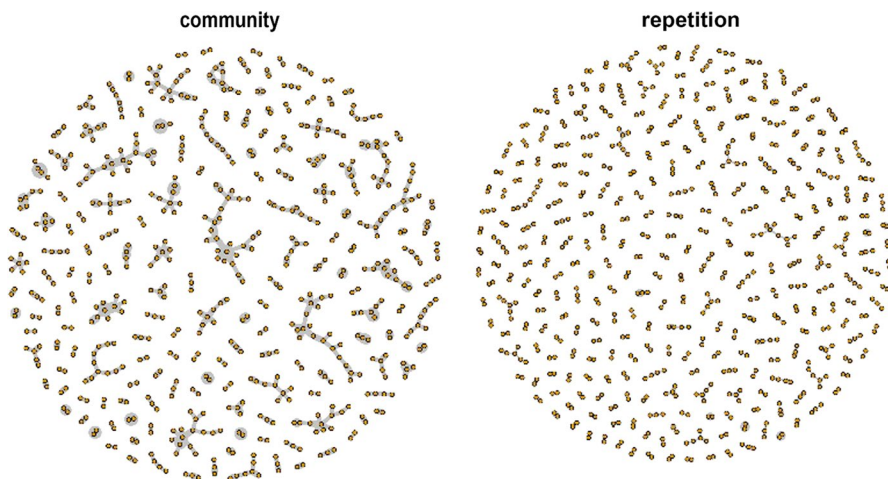


Fig. 2 Network formation along three periods (i.e., 3,000 interaction steps). Weak ties (i.e., ties representing a single interaction between agents) were removed to favor the graphical representation




Fig. 3 Network evolution along five periods (i.e., 5,000 steps). Different colors represent different interaction frequencies: 1 (gray), 2 (blue), and 3 or more (red)

In a targeted attack, experts no longer disseminate knowledge, thereby drastically reducing learning opportunities. This is an extremely realistic scenario, as discussed in the “Introduction.” To this end, we considered cases in which experts interrupted the dissemination of their superior knowledge after 10 and 50 periods, respectively, within simulations of 100 and 300 periods. After these turning points (i.e., 10 periods, 50 periods), dissemination was possible only among non-experts, as reported in the workflow.

In a random failure, interruption to knowledge dissemination affects both experts and non-experts (therefore, it is defined as random). As a case study, we considered a viral outbreak that temporarily deactivated certain nodes, due to infection. Even in this case, the interruption of interaction might affect knowledge diffusion, thereby increasing the level of inequality (Haelermans et al. 2022). The dynamics of this interruption may follow Block et al.’s (2020) S-I-R model, whereby q initial nodes are infected and then deactivated in the learning process. Considering the dynamics of contagion, we defined our model so that each subject had a probability of meeting other agents of $\pi_{\text{contact}} = 0.50$ and a probability of infecting those agents (if the subject was positive) of $\pi_{\text{infection}} = 0.80$. If infected, subjects could become infected after one period (i.e., 1,000 steps) and recover after four periods (i.e., 4,000 steps). When infected, the subject (node) would be deactivated and could no longer transmit knowledge. In this case, 50 periods were sufficient to track the virus spread.

3.5 Performance indicators

As in Morone and Taylor (2004), we considered proxies of knowledge and variance. Specifically, we analyzed:

- *Knowledge diffusion*, referring to the evolution of the aggregate number of knowledge tokens over time;
- The *Gini index*, referring to the overall inequality of knowledge in the system; and

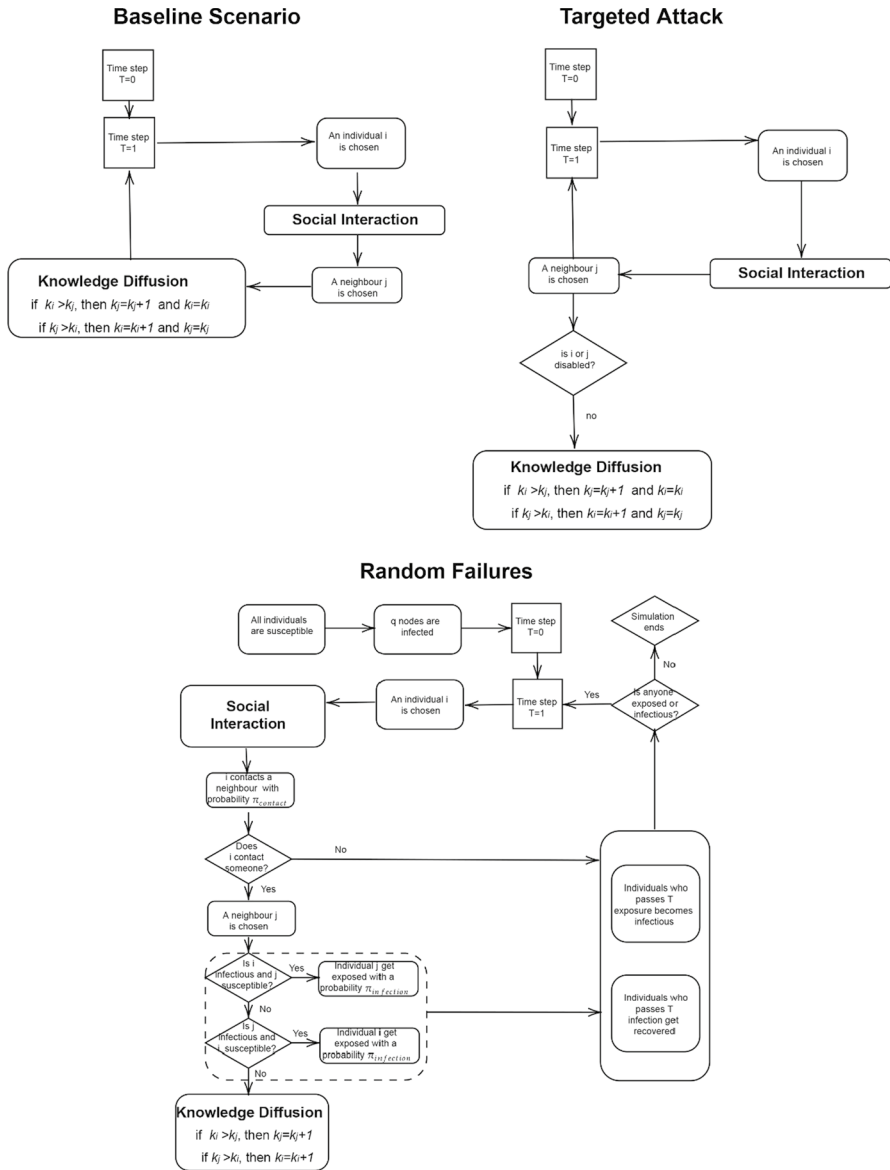


Fig. 4 Workflows for each scenario

- The *average knowledge gap*, as a robustness check of the Gini index, proxying the distance between the average knowledge of experts and the last 200 agents, in terms of the number of knowledge tokens held.

In line with the main research aim, the benchmark case was characterized by a suppression of inequality, indicated by a Gini index and knowledge gap close to 0.

However, as is observable in the results, this was difficult to achieve when the system was agitated.

4 Results

The following sections report the results, referring to the average of 30 simulations for each scenario.

4.1 Baseline scenario

Figure 5 presents the results for the baseline scenario. Interestingly, only few differences emerged among indicators, suggesting that the outcome was independent of the aggregation rule. This result was expected, since the scenario employed a static network (i.e., with the same potential interactions among subjects) and the different dynamics only reflected different (non-0) probabilities to select a given $i-j$ link. To

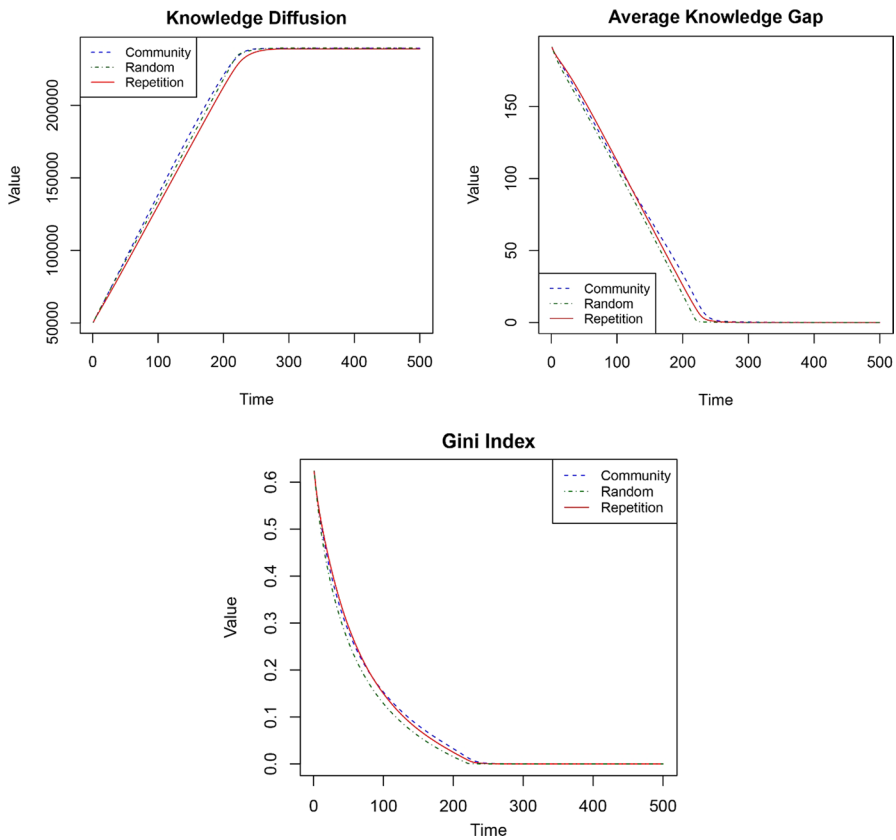


Fig. 5 Baseline results

clarify: in the random scenario, each link (within the local neighborhood defined in the static network at initialization) could be employed with equal probability, while in the repetition case, a subset of links (characterizing repeat interaction among the same dyads) had a higher probability and others had a lower probability of being employed. Consequently, in the random and community strategies, there was a higher probability that any potential link would be used; however, in the repetition case, many links had a low probability of being used, thereby greatly reducing the speed of dissemination and extending the time required to reach a steady state.

Considering the connected static network (with an average diameter of 3 and a mean degree of approximately 70), the differences between strategies were not defined by the long-run steady state achieved, but the time needed to reach knowledge saturation. In this respect, the community and the random strategies outperformed the repetition rule, requiring approximately 50 fewer periods to reach the steady state.

Considering the final configuration of knowledge diffusion, the irrelevance of social structure might suggest that social aspects are not relevant for policymaking, since agents will achieve the same results, independently of their manner of interaction. That is, the creation of scientific communities, industrial districts, or, more generally, organizational clusters seem irrelevant to the dissemination of knowledge. Additionally, a repetition structure, in which only a few contacts are informed by the same expert agent, also seems irrelevant. However, as the following results show, these considerations are myopic and do not consider the changing external circumstances that characterize the socioeconomic fabric.

4.2 Targeted attacks

This section reports the results of the models in which expert agents interrupted the communication process after 10 (Fig. 6) or 50 (Fig. 7) periods, as a robustness check. In this framework, the number of interactions was maintained as a constant. Thus, partner selection involved a trade-off between increasing the knowledge of one agent and reducing the learning opportunities of other nodes (i.e., a trade-off between the accumulation of knowledge (and human capital) by specific nodes at a lower level of diffusion and the distribution of knowledge to a larger slice of population at a higher level of diffusion but lower concentration).

Knowledge is crucial to preserve and maintain at a sophisticated level; however, this can lead to an unfavorable economic outcome characterized by a high concentration of knowledge and higher inequality, underlying the importance of knowledge diffusion. This trade-off is most evident under conditions of targeted attack, in which experts interrupt the diffusion of their knowledge. In these situations, further learning depends on: (i) the knowledge acquired by agents who were previously in contact with experts (i.e., before the attack) and (ii) the interaction between these agents and the remaining vertices. The first point requires intense and repeat interaction with experts, while the second requires diffused interaction.

The results show that the community strategy outperformed the other two strategies, in terms of both the aggregate level of knowledge diffusion and the

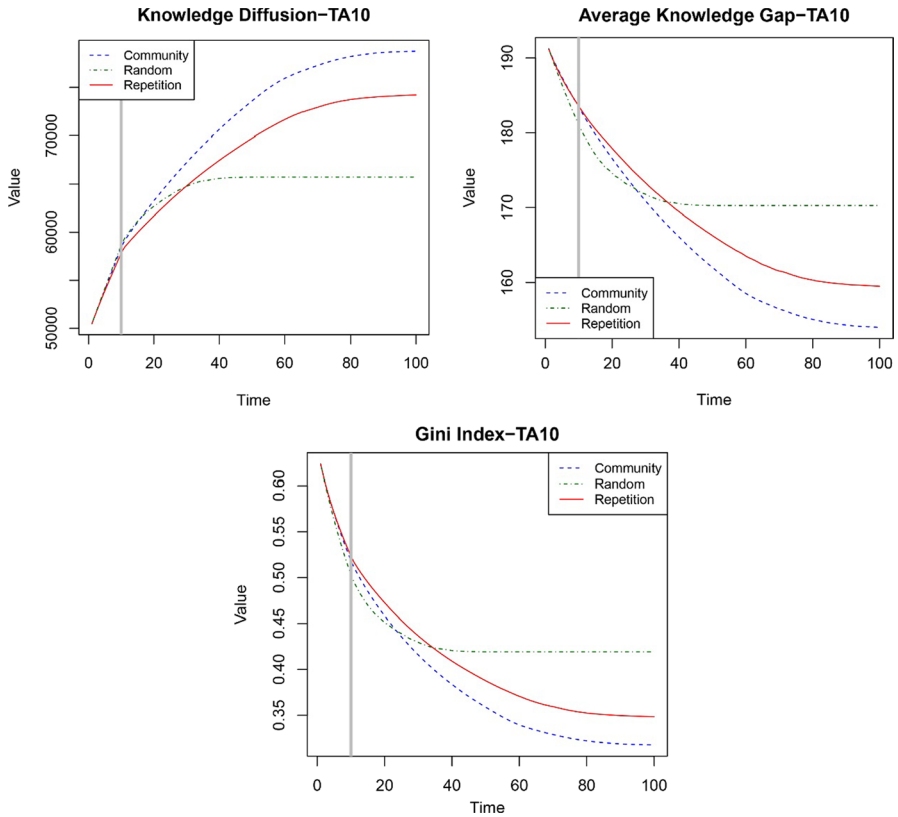


Fig. 6 Network under targeted attack after 10 periods

reduction of inequality. After 10 (50) periods, in the community strategy, experts diffused sufficient knowledge and continued to propagate, establishing more learning opportunities after excluding the most contacted nodes. In the repetition strategy, the diffusion rate was lower and the concentration of knowledge higher, but the exclusion of experts did not allow nodes that had acquired a high level of knowledge to continue in the diffusion process. Therefore, there were fewer learning opportunities. In the random scenario, the network was open and knowledge was diffuse, but also dispersed. To wit, after the removal of experts, the level of knowledge among non-experts was low, resulting in no further learning opportunities.

Thus, the community strategy reflected a compromise between concentrating knowledge in a few agents and diffusing knowledge throughout the network. Bonding capital (i.e., strong ties, or intense interactions within communities) allowed for the “creation” of experts with higher knowledge, while bridging capital (i.e., weak ties, or interactions with the rest of the neighborhood) created knowledge spillover, thereby successfully increasing learning opportunities. Interestingly, when there was no organization in the diffusion of knowledge (i.e.,

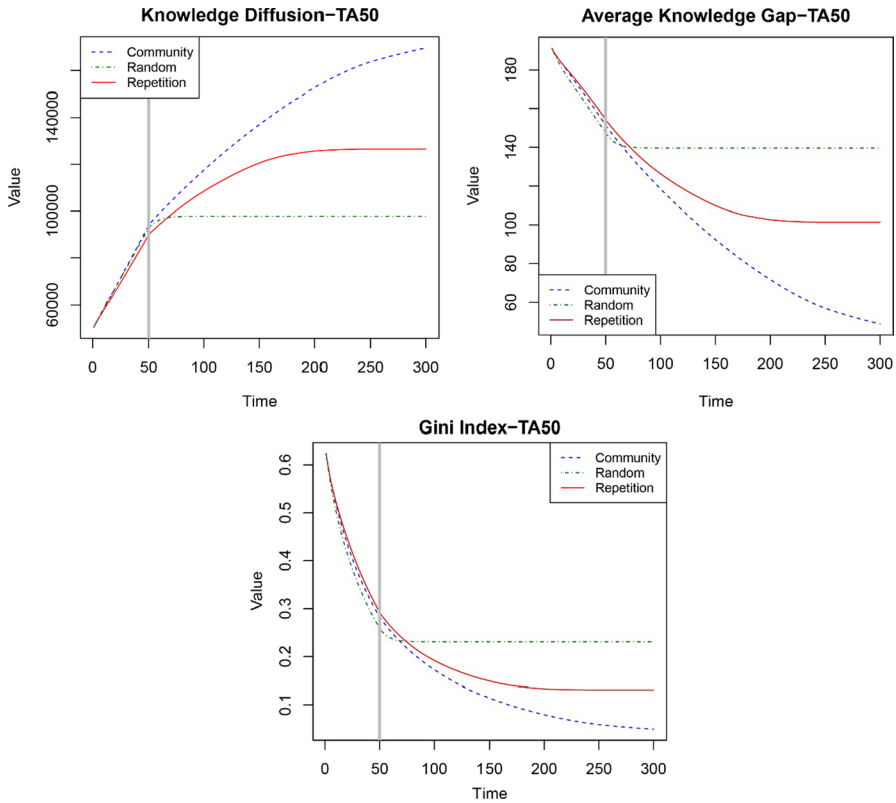


Fig. 7 Network under targeted attack after 50 periods

in the random strategy), knowledge was diffused but dispersed, and the emerging network was not resilient to external structural change. Thus, clusters, scientific communities, and, more generally, melting pots, may support both the preservation of knowledge and knowledge dissemination, even in the context of drastic change in the socioeconomic environment. Figure 8 summarizes this aspect, comparing the knowledge distribution across the three strategies at the end of the resilience test (TA50 is taken as example), also including the starting point, that is the initial distribution of knowledge (Fig. 1). Similarly to the first figure showed, on the x-axis, the agents are sorted in descending order according to their knowledge level. It is evident that the community strategy exhibits a high level of knowledge and low level of inequalities between experts and non-experts.

4.3 Random failure

Even the results under conditions of random failure confirmed the superior performance of the community strategy. Differently from the previous scenario of a targeted attack, in this case, we did not expect long-run differences, given the

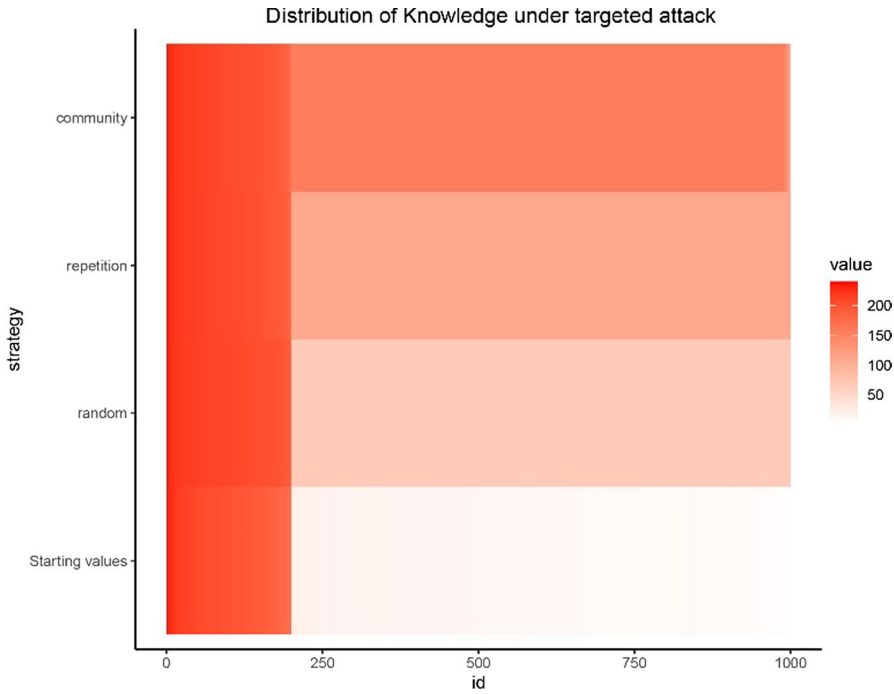


Fig. 8 Heatmap reporting the final distribution of knowledge across strategies (for TA-50, the most representative case with higher differences). On the x-axis agents are sorted in descending order according to their level of knowledge (filled red area)

short-term nature of the shock (i.e., agents were only randomly and temporarily disabled, and interactions were fully restored after a small number of periods). Indeed, we did not observe long-run changes in inequality measures. Figure 9 reports the growth in knowledge diffusion given the short-term deactivation of nodes but no

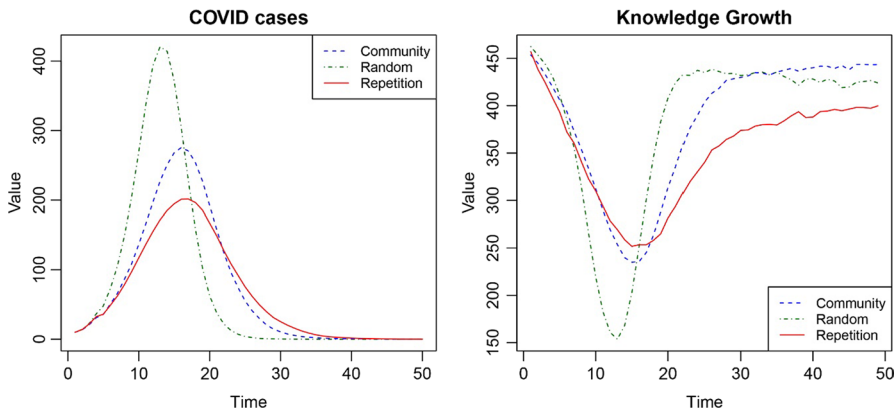


Fig. 9 Network under conditions of random failure

long-term change in the aggregate. Although the number of temporarily inactive nodes (e.g., COVID-infected cases) was slightly higher in the community case than the repetition case, the organizational structure experienced a faster recovery, showing greater growth rates after the initial decline. However, no significant differences were observed in the change in inequality. (Therefore, these results are omitted.)

The random scenario outlined the sharpest decline in growth, due to the high number of inactive nodes. However, the community scenario preserved networks of contact that could mitigate the excessive spread of the virus (given the high number of intra-group contacts), while at the same time preserving the interactive learning process.

5 Discussion

In the face of numerous global challenges to humanity, the concept of resilience is gaining momentum. Different from individualistic perspectives, resilience may be understood as a social process linked to the social context, characterized by human interaction in varying environments (Revilla et al. 2018).

The structural components of social networks may contribute significantly to defining their economic output and maintaining the dissemination of behavioral and informational knowledge, even during conditions of shock. In particular, organized networks seem to function better than networks that rely on random social interactions. Furthermore, societies that foster both network openness and stable micro-organizations tend to enjoy better knowledge dissemination, even during negative shocks.

The present analysis showed that structured networks that balance inter-group strong ties (favoring the creation of communities) and weak ties (favoring knowledge spillover across communities) perform the best in situations of adversity. Put differently, this organizational scheme seems to guarantee the highest aggregate level of knowledge and lowest level of inequality, and thereby the greatest resilience.

The analysis modeled two types of negative shocks that might alter and undermine communication: a targeted attack, interrupting the knowledge dissemination of experts, and a random failure, interrupting the communication of both experts and non-experts. The first case could emerge from the desire of a specific group to maintain a high level of social inequality. As reported in the “Introduction,” high-tech firms can collude to prevent knowledge diffusion outside their cluster, in order to maintain a competitive advantage (Bacchiega et al. 2010); experts might also be interested in preserving their economic advantage in different contexts, keeping (at least) constant the gap between the rich and the poor. Random failures, in which the randomness refers to the types of nodes that are no longer able to process or disseminate knowledge, are also evident in the real world, with the recent COVID-19 pandemic representing but one example (Haelermans et al. 2022).

Inequalities of knowledge reflect inequalities of opportunity, and subsequently economic inequalities. As largely agreed, the perpetuation of inequalities is dangerous for the stability of any socioeconomic system (Stiglitz 2016). The present analysis attempted to solve the efficiency versus equity trade-off (Cowan

and Jonard 2004) faced by policymakers in the area of regional technology and knowledge. The findings revealed a network structure resulting in a high average level of knowledge (i.e., high efficiency) and low heterogeneity (i.e., high equity) among agents.

These results underline the importance of examining social aggregation rules in the context of communication, and the importance of preserving communities (e.g., scientific communities, consortia, general melting pots) and organizational clusters to ensure the inheritance, preservation, and transmission of knowledge. The study contributes a social perspective to the literature that might guide individuals, firms, and social agents, more generally, to define social trajectories that not only preserve efficiency and equity, but also build resilience.

From a policy perspective, the creation of communities and melting pots may mitigate the impact of adverse phenomena, promoting higher stability and resilience within the wider social network. Therefore, we can conclude that the co-existence of both weak and strong ties is necessary in any social organization to ensure the successful acquisition and exchange of knowledge through collaborative interaction.

Appendix A

The following reports the computational steps applied for the definition of social interactions.

Step 1: Agent distribution and attributes. N agents allocated in a 2D space. Each agent i has different attributes:

- **geographical collocation:** x – y coordinates;
- **group ID:** a random ID number assigning agent i to a group, whereby two agents i and j sharing the same group ID are members of the same group; and
- **homophily attribute:** a random scalar a_i ranging from 0–100, with the closest value represented by the value between the vertexes i and j and the higher value representing their homophily.

Step 2: Definition of the static network. For each agent i , a potential list of contacts is generated, considering:

- **geographical proximity:** contacts sorted by Euclidean distance and randomly generated in sets ranging from 4–12 (assuming a probability of $d_{\text{geo}}=0.3$);
- **group contacts:** groups composed of an average of eight members, with each actor forming ties at random with other members (assuming an agent density of $g_{\text{groups}}=0.6$);
- **homophilistic proximity:** the normalized absolute distance considering the scalar a_i and a random number of ties with the closest alters (ranging from 4–12); and
- **random contacts:** each actor having a random partner in their neighborhood with a probability of $z=0.5$.

Step 3: Definition of the dynamic network statistics. Considering N_i as the neighborhood of agent i , within each step, a randomly selected agent chooses a partner according to a probability model p . Each agent j in the neighborhood of i has a probability of being selected that depends on the behavioral strategy adopted:

- **random strategy:** each contact j has the same probability of being selected;
- **community strategy:** each contact j has a higher probability of being selected, depending on the following statistic $s(i, j)$:

$$s_{\text{community}}(i, j) = \sum_{h=1}^N x_{i,h} x_{j,h}$$

where h is a third alter different from i and j , and $x_{i,h} x_{j,h}$ is equal to 1 whenever alter h is in the contact list of both i and j , and otherwise is equal to 0; and **repetition strategy:** the probability of selecting a partner is driven by the following statistic $s(i, j)$:

$$s_{\text{repetition}}(i, j) = L_i(j, t)$$

where t is the time step of the simulation and $L_i(j, t)$ is the number of times i met j previously, considering a time window of $\lambda = 3$ (i.e., the previous three interactions for each agent).

Step 4. Definition of the dynamic network probability model. Given these features and following a multinomial choice model, each statistic is weighted by a parameter α and calibrated to make the models comparable (see Block et al. 2020). Therefore, the probability of selecting a contact j is given by:

$$p(i \rightarrow j | \alpha) = \frac{\exp(\alpha * s(i, j))}{\sum_{j \in N} \exp(\alpha * s(i, j))}$$

Such that, given two potential partners of i , j_1 and j_2 , when the statistic increases by 1 unit ($s(i, j_2) = s(i, j_1) + 1$), the marginal increase in the probability of selecting j_2 is:

$$\log \frac{p(i \rightarrow j_2 | \alpha)}{p(i \rightarrow j_1 | \alpha)} = \alpha$$

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Alatas V, Banerjee A, Chandrasekhar AG, Hanna R, Olken BA (2016) Network structure and the aggregation of information: theory and evidence from Indonesia. *American Econom Rev* 106(7):1663–1704

- Ally AF, Zhang N (2018) Effects of rewiring strategies on information spreading in complex dynamic networks. *Commun Nonlinear Sci Numer Simul* 57:97–110
- Bacchiaga E, Lambertini L, Mantovani A (2010) R&D-hindering collusion. *BE J Econom Anal Policy*. <https://doi.org/10.2202/1935-1682.2157>
- Bala V, Goyal S (2001) Conformism and diversity under social learning. *Econ Theor* 17:101–120
- Block P, Hoffman M, Raabe IJ, Dowd JB, Rahal C, Kashyap R, Mills MC (2020) Social network-based distancing strategies to flatten the COVID-19 curve in a post-lockdown world. *Nat Hum Behav* 4(6):588–596
- Breetzke J, Wild EM (2022) Social connections at work and mental health during the first wave of the COVID-19 pandemic: evidence from employees in Germany. *PLoS ONE* 17(6):e0264602
- Clementi F, Gallegati M (2005) Pareto's law of income distribution: evidence for Germany, the United Kingdom, and the United States. Springer. In *Econophysics of wealth distributions*
- Cowan R, Jonard N (2003) The dynamics of collective invention. *J Econ Behav Organ* 52(4):513–532
- Cowan R, Jonard N (2004) Network structure and the diffusion of knowledge. *J Econ Dyn Control* 28(8):1557–1575
- De Domenico M, Solé-Ribalta A, Gómez S, Arenas A (2014) Navigability of interconnected networks under random failures. *Proc Natl Acad Sci* 111(23):8351–8356
- Frank D (ed) *The new economic sociology: a reader* Princeton. Princeton University Press, NJ
- Dong G, Gao J, Du R, Tian L, Stanley HE, Havlin S (2013) Robustness of network of networks under targeted attack. *Phys Rev E* 87(5):052804
- Granovetter MS (1973) The strength of weak ties. *Am J Sociol* 78(6):1360–1380
- Greenan CC (2015) Diffusion of innovations in dynamic networks. *J Royal Stat Soc A (stat Soc)* 178:147–166
- Günther J, Meissner D (2017) Clusters as innovative melting pots? The meaning of cluster management for knowledge diffusion in clusters. *J Knowl Econ* 8(2):499–512
- Haelermans C, Korthals R, Jacobs M, de Leeuw S, Vermeulen S, van Vugt L, de Wolf I (2022) Sharp increase in inequality in education in times of the COVID-19 pandemic. *PLoS ONE* 17(2):e0261114
- Jialu S, Zhiqiang M, Binxin Z, Haoyang X, Fredrick Oteng A, Hu W (2021) Collaborative innovation network, knowledge base, and technological innovation performance-thinking in response to COVID-19. *Front Psychol* 12:648276
- Kim H, Park Y (2009) Structural effects of R&D collaboration network on knowledge diffusion performance. *Expert Syst Appl* 36(5):8986–8992
- Kirman A (2010) *Complex economics: individual and collective rationality*. Routledge
- Lambiotte R, Panzarasa P (2009) Communities, knowledge creation, and information diffusion. *J Informet* 3(3):180–190
- Latif MA, Naveed M, Zaidi F (2013). Resilience of social networks under different attack strategies. In: *International Conference on Social Informatics*. Springer.
- Li W, Li Y, Tan Y, Cao Y, Chen C, Cai Y, Pecht M (2019) Maximizing network resilience against malicious attacks. *Sci Rep* 9(1):1–9
- Lin M, Li N (2010) Scale-free network provides an optimal pattern for knowledge transfer. *Physica A Stat Mech Appl* 389(3):473–480
- Liu C, Zhang ZK (2014) Information spreading on dynamic social networks. *Commun Nonlinear Sci Numer Simul* 19(4):896–904
- Malmberg A, Power D (2005) (How) do (firms in) clusters create knowledge? *Ind Innov* 12(4):409–431
- May RM, Lloyd AL (2001) Infection dynamics on scale-free networks. *Phys Rev E* 64(6):066112
- Midgley DF, Morrison PD, Roberts JH (1992) The effect of network structure in industrial diffusion processes. *Res Policy* 21(6):533–552
- Morone P, Taylor R (2004) Knowledge diffusion dynamics and network properties of face-to-face interactions. *J Evol Econ* 14(3):327–351
- Nahapiet J, Ghoshal S (1998) Social capital, intellectual capital, and the organizational advantage. *Acad Manag Rev* 23(2):242–266
- Nelson RE (1989) The strength of strong ties: social networks and intergroup conflict in organizations. *Acad Manag J* 32(2):377–401
- Putnam R (2001) Social capital: measurement and consequences. *Can J Policy Res* 2(1):41–51
- Rappa MA, Debackere K (1992) Technological communities and the diffusion of knowledge. *R&D Manag* 22(3):209–220
- Revilla JC, Martín P, de Castro C (2018) The reconstruction of resilience as a social and collective phenomenon: poverty and coping capacity during the economic crisis. *Eur Soc* 20(1):89–110

- Rogers EM (2003) Diffusion of innovations, 5th edn. Free Press, Newyork
- Schuller T, Theisens H (2010). Networks and communities of knowledge. *Int Encyclop Edu*
- Snijders TA (2001) The statistical evaluation of social network dynamics. *Sociol Methodol* 31(1):361–395
- Snijders TA, Van de Bunt GG, Steglich CE (2010) Introduction to stochastic actor-based models for network dynamics. *Social Netw* 32(1):44–60
- Steglich C, Snijders TA, Pearson M (2010) Dynamic networks and behavior: separating selection from influence. *Soc Methodol* 40(1):329–393
- Stiglitz JE (2016). Inequality and economic growth. *Political Quarterly*
- Storck J, Hill PA (2009). Knowledge diffusion through strategic communities, In: knowledge and communities, Routledge.
- Strang D, Tuma NB (1993) Spatial and temporal heterogeneity in diffusion. *Am J Sociol* 99(3):614–639
- Sumpter DJ (2010) Collective animal behavior. Princeton Univ Press, In *Collective animal behavior*
- Tang F, Mu J, MacLachlan DL (2010) Disseminative capacity, organizational structure and knowledge transfer. *Expert Syst Appl* 37(2):1586–1593
- Todo Y, Matous P, Inoue H (2016) The strength of long ties and the weakness of strong ties: knowledge diffusion through supply chain networks. *Res Policy* 45(9):1890–1906
- Woolcock M (1998) Social capital and economic development: toward a theoretical synthesis and policy framework. *Theory Soc* 27(2):151–208
- Yang GY, Hu ZL, Liu JG (2015) Knowledge diffusion in the collaboration hypernetwork. *Physica A* 419:429–436
- Zhang H, Wu W, Zhao L (2016) A study of knowledge supernetworks and network robustness in different business incubators. *Physica A Stat Mech Appl* 447:545–560
- Zhao J, Bai A, Xi X, Huang Y, Wang S (2020) Impacts of malicious attacks on robustness of knowledge networks: a multi-agent-based simulation. *J Knowl Manag* 24(5):1079–1106
- Zheng W, Pan H, Sun C (2019) A friendship-based altruistic incentive knowledge diffusion model in social networks. *Inf Sci* 491:138–150

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.