



Does Spending More Always Ensure Higher Cooperation? An Analysis of Institutional Incentives on Heterogeneous Networks

Theodor Cimpanu¹ · Francisco C. Santos² · The Anh Han³

Accepted: 14 March 2023
© The Author(s) 2023

Abstract

Humans have developed considerable machinery used at scale to create policies and to distribute incentives, yet we are forever seeking ways in which to improve upon these, our institutions. Especially when funding is limited, it is imperative to optimise spending without sacrificing positive outcomes, a challenge which has often been approached within several areas of social, life and engineering sciences. These studies often neglect the availability of information, cost restraints or the underlying complex network structures, which define real-world populations. Here, we have extended these models, including the aforementioned concerns, but also tested the robustness of their findings to stochastic social learning paradigms. Akin to real-world decisions on how best to distribute endowments, we study several incentive schemes, which consider information about the overall population, local neighbourhoods or the level of influence which a cooperative node has in the network, selectively rewarding cooperative behaviour if certain criteria are met. Following a transition towards a more realistic network setting and stochastic behavioural update rule, we found that carelessly promoting cooperators can often lead to their downfall in socially diverse settings. These emergent cyclic patterns not only damage cooperation, but also decimate the budgets of external investors. Our findings highlight the complexity of designing effective and cogent investment policies in socially diverse populations.

Keywords Evolutionary game theory · Evolution of cooperation · Cost efficiency · Incentives · Scale-free networks · Prisoners' dilemma

This article is part of the topical collection "Evolutionary Games and Applications" edited by Christian Hilbe, Maria Kleshnina and Kateřina Staňková.

✉ Theodor Cimpanu
tic1@st-andrews.ac.uk

✉ The Anh Han
T.Han@tees.ac.uk

¹ School of Mathematics and Statistics, University of St Andrews, St Andrews, UK

² INESC-ID and Instituto Superior Tecnico, Universidade de Lisboa, Lisbon, Portugal

³ School Computing, Engineering and Digital Technologies, Teesside University, Middlesbrough, UK

1 Introduction

The design of mechanisms that encourage pro-social behaviours in populations of self-regarding agents is recognised as a major theoretical challenge within several areas of social, life and engineering sciences. It is ubiquitous in real-world situations, not least ecosystems, human organisations, technological innovations and social networks [3, 14, 23, 43, 47, 52]. In this context, cooperation is typically assumed to emerge from the combined actions of individuals within the system. However, in many scenarios, such behaviours are advocated and promoted by an external party, which is not part of the system, calling for a new set of heuristics capable of *engineering* a desired collective behaviour in a self-organised complex system [38]. Among these heuristics, several have been identified as capable of promoting desired behaviours at a minimal cost [10, 11, 13, 18, 19, 22, 26, 59]. However, these studies neglect the diversified nature of contexts and social structures which define real-world populations. Here, we analyse the impact of diversity by means of scale-free interaction networks with dissimilar levels of clustering [6, 46] and test various interference mechanisms using simulations of agents facing a cooperative dilemma.

For instance, if one considers a near future, where hybrid societies comprising humans and machines shall prevail, it is important to identify the most effective incentives to be included to leveraging cooperation in such hybrid collectives [3, 14, 24, 37]. In a different context, let us consider a wildlife management organisation (e.g. the WWF) that aims to maintain a desired level of biodiversity in a particular region. In order to do that, the organisation, not being part of the region's ecosystem, has to decide whether to modify the current population of some species, and if so, then when and in what degree to *interfere* in the ecosystem (i.e. to modify the composition of the population) [29]. Since a more impactful intervention typically implies larger costs in terms of human resources and equipment, the organisation has to achieve a balance between cogent wildlife management and a low total investment cost. Moreover, due to the evolutionary dynamics of the ecosystem (e.g. frequency and structure dependence) [28, 30, 47], undesired behaviours can reoccur over time, for example when the interference was not sufficiently strong in the past. Given this, the decision-maker also has to take into account the fact that it will have to repeatedly interfere in the ecosystem in order to sustain levels of biodiversity over time. That is, they must find an efficient interference mechanism that leads to their desired goals, while also keeping in mind potential budget concerns.

Moreover, real-world networks of individuals, such as social networks and networks of collaboration, are inherently *heterogeneous* [6]. Moreover, in the context of evolutionary game theory (EGT), scale-free networks imply more than the underlying interaction structure. Heterogeneous graphs can portray social diversity [48] and the inherent inequality that exists between individuals. Individuals vary in influence and accumulated wealth, and this diversity has been shown to play a key role in the evolution of cooperative behaviours [42, 47], enhancing the resilience of cooperation by inducing cooperative agents to create assortative clusters, where they reciprocate cooperation. Notwithstanding, several recent works have also discovered that heterogeneity could play a negative role in the evolution of cooperation, in the context of weak selection [2, 31, 32, 55]. These findings lend particular importance to carefully considering network structures when designing incentive schemes.

Specifically, we consider populations of individuals distributed in a scale-free network, who interact with their neighbours via the one-shot prisoner's dilemma (PD), where uncooperative behaviour is preferred over cooperation [40, 44, 47, 51]. As an outside decision-maker, we aim to promote cooperation by interfering in the system, rewarding particular agents in

the population at specific moments. The research question here is to identify when and how much to invest (in individuals distributed in a network) at each time step, in order to achieve cooperation within the system such that the total cost of interference is minimised, taking into account the fact that individuals might have different levels of social connectivity. For instance, we might wonder whether it is sufficient to focus the investment only on highly connected cooperators, as they are more influential. Would targeting influencers reduce overall costs? Do we need to take into account a neighbourhood's cooperativeness level, which was shown to play an important role in square lattice networks [22]? Also, when local information is not available and only global statistics can be used in the decision-making process, how different are the results in heterogeneous networks, in comparison with regular graphs (i.e. homogeneous networks)?

To answer these questions, this paper will systematically investigate different general classes or approaches of interference mechanisms, which are based on (i) the global population statistics such as its current composition, (ii) a node's degree centrality in the network and (iii) the neighbourhood properties, such as local cooperativeness level.

Our results show that interference on scale-free networks is not trivial. In particular, we highlight that the inconsiderate distribution of incentives can lead to the exploitation of cooperators. We present which mechanisms are more efficient at fostering cooperation, arguing that social diversity and the network's clustering coefficient both play a key role in the choice of interference mechanisms available to institutions wishing to promote cooperation.

2 Model and Methods

2.1 Prisoner's Dilemma on Scale-Free Networks

We consider a population of agents on scale-free networks of contacts (SF NoCs)—a widely adopted heterogeneous population structure in population dynamics and evolutionary games. We focus our analysis on the efficiency of various interference mechanisms in spatial settings, adopting an agent-based model directly comparable with the setup of recent laboratory experiments on cooperation [45].

One of the main targets of this work is therefore to measure the impact of network properties and structural heterogeneity in the evolutionary dynamics of interference behaviours. To measure these effects, we will make use of two types of scale-free networks [6, 16, 34], generated through two growing network models: the *Barabási and Albert* (BA) model [6] and the *Dorogovtsev–Mendes–Samukhin* (DMS) model [17].

The Barabási and Albert (BA) model [6] is one of the most famous models used in the study of highly heterogeneous, complex networks. The main features of the BA model are that it follows a *preferential attachment* rule and has a small clustering coefficient and a typical *power-law degree distribution*. In order to explain preferential attachment, let us describe the construction of a BA network. Starting from a small set of m_0 interconnected nodes, each new node selects and creates a link with m older nodes according to a probability proportional to their degree. The procedure stops when the required network size of N is reached. This will produce a network characterised by a power-law distribution,

$$p_k \sim k^{-\gamma}, \quad (1)$$

where the exponent γ is its *degree exponent* [5]. There is a high degree correlation between nodes, and the degree distribution is typically skewed with a long tail. There are few hubs

in the network that attract an increasing number of new nodes which attach as the network grows (in a typical “*rich-get-richer*” scenario). The power-law distribution exhibited by BA networks resembles the heterogeneity present in many real-world networks. However, they are also defined by low clustering coefficients, which means they cannot always be used to approximate realistic settings [54].

To build heterogeneous networks with a large clustering coefficient, [17] have proposed the eponymous DMS model. This model follows a similar method of construction as the BA model, is also exemplary of the preferential attachment rules and follows a power-law degree distribution. Crucially, each new node connects with the two extremities of m ($m \geq 2$) randomly chosen edges, instead, therefore forming characteristic triangular motifs whenever a new node is added to the network. Since the number of edges arriving to any node reflects its degree, the probability of attaching the new node to an old node is proportional to its degree and preferential attachment is recovered. The degree distribution is therefore the same as the one of a BA model, and the degree–degree correlations are also equal [15]. However, the clustering coefficient is large, and more accurately mimics many realistic social networks [7, 54]. Moreover, we select an initial number of nodes $m_0 = 2$, with two additional edges being created at every time step of network generation. This produces networks of average connectivity $z = 4$, serving as a direct comparison between this work and other studies performed on structured populations [22].

Initially each agent in a population of size N is designated as either a cooperator (C) or defector (D) with equal probability. Agents’ interaction is modelled using the one-shot prisoner’s dilemma game, where mutual cooperation (mutual defection) yields the reward R (penalty P) and unilateral cooperation gives the cooperator the sucker’s pay-off S and the defector the temptation T . As a popular interaction model of structured populations [36], we adopt the following scaled pay-off matrix of the PD (for row player):

$$\begin{matrix} & C & D \\ C & \begin{pmatrix} 1 & 0 \\ b & 0 \end{pmatrix} \\ D & \end{matrix}$$

with b ($1 < b \leq 2$) representing the temptation to defect. We adopt this weak version of the prisoner’s dilemma in spite of cooperation prevalence shown in previous works on scale-free networks [48], so as to have a direct comparison with studies on the effects of rewarding mechanisms in different types of networks [22].

At each time step or generation, each agent plays the PD with its immediate neighbours. The score for each agent is the sum of the pay-offs in these encounters. Before the start of the next generation, the conditions of interference are checked for each agent, and if they qualify, the external decision-maker increases their pay-off. Multiple mechanisms (i.e. multiple conditions) can be active at once, but the individual investment cannot be applied more than once; the schemes determine the eligibility for investment.

At the start of the next generation, each agent’s strategy is updated using one of two social learning paradigms—a *deterministic* or a *stochastic* rule. Using a deterministic update rule, each agent will choose to imitate the strategy of its highest scored neighbour [36, 56]. In the stochastic case, instead of copying the highest scored neighbour, at the end of each generation an agent A with score f_A chooses to copy the strategy of a randomly selected neighbour agent B with fitness f_B , using a pairwise comparison rule, with an imitation probability given by the Fermi–Dirac function [57]:

$$(1 + e^{(f_A - f_B)/K})^{-1},$$

where K denotes the amplitude of noise in the imitation process [56]. We set $K = 0.1$ in our simulations, a value usually adopted in previous works [11, 41, 56]. Our analysis will be based on this standard evolutionary process in order to focus on understanding the cost efficiency of different interference mechanisms.

We simulate this evolutionary process until a stationary state or a cyclic pattern is reached. The simulations converge quickly in the case of deterministic update, with the exception of some cyclic patterns which never reach a stationary state. Because this work studies cost-effective intervention, these rarely occurring patterns which inherently invite very large total costs are escaped early by running simulations for 75 generations (deterministic update) and 500 generations (stochastic update), at which point the accumulated costs are excessive enough for this mechanism to not be of interest. The difference in the final number of generations accounts for the slower convergence time associated with stochastic dynamics. Moreover, the results are averaged for the last 25 generations of the simulations for a clear and fair comparison (e.g. due to cyclic patterns). In order to improve accuracy related to the randomness of network topology in scale-free networks, each set of parameter values is ran on 10 different pre-seeded graphs for both types of SF NOCs. Furthermore, the results for each combination of network and parameter values are obtained from averaging 30 independent realisations. It is important to note that the distribution of cooperators and defectors on the network is different for every realisation.

Note that we do not consider mutations or random explorations when employing a deterministic update rule. Thus, whenever the population reaches a homogeneous state (i.e. when the population consists of 100% of agents adopting the same strategy), it will remain in that state regardless of interference. Hence, whenever detecting such a state, no further interference will be made. Errors can sometimes occur under the presence of stochastic imitation; thus, we never pre-emptively pause these simulations. Given the difference in convergence time, network size and stopping conditions, we do not directly compare the total costs between these two paradigms.

2.2 Cost-Efficient Interference in Networks

We aim to study how one can efficiently interfere in spatially heterogeneous populations to achieve high levels of cooperation while minimising the cost of interference. An investment decision consists of a cost $\theta > 0$ to the external decision-making agent/investor, and this value θ is added as surplus to the pay-off of each suitable candidate. In order to determine cost efficiency, we vary θ for each proposed interference strategy, measuring the total accumulated costs to the investor. Thus, the most efficient interference schemes will be the ones with the lowest relative total cost.

Moreover, in line with previous works on network interference [10, 22, 26], we compare global interference strategies where investments are triggered based on network-wide information, local neighbourhood information and, lastly, node centrality information.

In the *population-based (POP)* approach, a decision to invest in desirable behaviours is based on the current composition of the population. We denote x_c the fraction of individuals in the population adopting cooperative behaviour. Namely, an investment is made if x_c is at most equal to a threshold p_c (i.e. when $x_c \leq p_c$), for $0 \leq p_c \leq 1$. They do not invest otherwise (i.e. $x_c > p_c$). The value p_c describes how rare the desirable behaviours should be to trigger external support.

In the *neighbourhood-based (NEB)* approach, committing an abuse of notation, a decision to invest is based on the fraction x_c of neighbours of a focal individual with the desirable

behaviours, calculated at the local level. Investment happens if x_c is at most equal to a threshold n_c (i.e. when $x_c \leq n_c$), for $0 \leq n_c \leq 1$; otherwise, no investment is made.

Given the deliberate inequality between nodes, we need a measurement to distinguish important/influential nodes from less important ones. The degree centrality is the oldest measure of importance or influence ever used in network science [9]. It denotes the number of neighbours of the node i ; namely, it measures the number of edges of node i . By definition, degree centrality is normalised using the total number of nodes, or the maximal possible degree, $n - 1$, to obtain a number between 0 and 1. Despite its simple definition, degree centrality is often a highly effective measure of the influence or importance of a node, since people with more connections tend to be more influential in a social network [8, 35]. The reason why we define degree centrality by using the previous normalised definition, and not simply the degree, is that it allows us to compare two nodes that belong to different networks regardless of network size [50]. We will refer to this examination as the *node influence-based (NI)* approach. Degree centrality, denoted by NI-deg or x_i^{deg} , is defined as follows:

$$x_i^{deg} = deg_i = \frac{k_i}{n - 1}, \quad (2)$$

where k_i is the degree of the node i and $n - 1$ is the total number of nodes. The degree k_i of a node i is given by: $k_i = \sum_{j=1}^n A_{ij}$, where A is the adjacency matrix of a finite graph, populated with pairs of vertices which are adjacent (i.e. connected). The decision-maker invests in a cooperator node C when the value of its degree centrality is above a threshold of influence c_I , for $0 \leq c_I \leq 1$. Otherwise, i.e. when $0 \leq x_i < c_I$, no investment is made. The value c_I describes how influential a cooperator node should be to trigger an investment into its survival.

For the POP and NEB schemes, the threshold signifies an increase in the number of nodes that satisfy the requirements for investment. In other words, a threshold of 1 means always investing in all nodes which follow the desired strategy. Conversely, a lower threshold implies a more careful approach to investment, whereby the exogenous agent is stricter in their selection of suitable candidates. The opposite is true for NI, as a value of 1 implies only the most connected individual(s) is eligible for investment, whereas a value of 0 means investing in every cooperative agent.

Interestingly, we posit that these mechanisms require different levels of information, which may or may not be readily available in the given network. In some cases, such as social networks, the connectivity (i.e. the number of friends) of a node is virtually free information which requires no effort on the part of the external decision-maker to discern. On the other hand, neighbourhood-based approaches inherently require more information about the population and the level of cooperativeness in different parts of the network. Thus, POP is a broad mechanism which only requires knowledge about overall cooperativeness, but NEB invites complex information gathering, in order to determine the cooperativeness in each neighbourhood. Combining NI with NEB does not require any additional observation than NEB by itself. Our study of neighbourhood-based interference does not directly take into account the cost of gathering information, it is a comparison between perceived gains in cooperation and the associated per-individual cost of interference set out in the interference mechanisms. Our discussion will naturally present these subtle differences in the hierarchy of information gathering, as they signal hidden costs for some application domains.

3 Results

In contrast to the study on square lattice networks [22], we found that performing cost-effective interventions on SF NOCs exhibits complex patterns and presents multiple concerns. In structured populations, more detailed observations resulted in effective interventions with improved outcomes. On the other hand, more knowledge about the population in SF NOCs simply reduces the risk of interfering to the detriment of cooperators. In other words, interfering in SF NOCs without adequate knowledge should be approached cautiously or it could act to the benefit of defectors. This issue is prevalent in SF networks with low clustering (BA model), but also sees some representation in highly clustered (DMS) networks if stochastic dynamics are taken into account.

Successfully investing in BA populations broadly requires heavy-handed investment and large individual endowments (often orders of magnitude higher than similar mechanisms performed on square lattice populations) or a blanketing mechanism that targets all or almost all cooperators, even those which are not necessarily in danger of converting to defection. Converging to 100% C is very difficult unless both of these conditions are met and this introduces multiple concerns in the role of an exogenous interfering party. We avoid focusing on solutions where the per-generation cost is excessive, as it is unlikely for any institution to be able to produce unrealistically high endowments, as required by these heterogeneous networks. Instead we focus on effective intervention with manageable amounts of per-generation cost. In the following subsections, we structure our results based upon the most important findings and provide relevant references to each studied investment scheme where appropriate. Initially, we will present the results for the deterministic update scheme and then the stochastic update, pointing out any difference between the two. All the main findings are robust irrespective of the social learning paradigm employed.

3.1 Careless Rewarding Leads to the Exploitation of Cooperators

In direct contrast to previous findings for positive incentives [10, 22, 26], an external decision-maker should only interfere in scale-free networks with great care, as investing indiscriminately can lead to the detriment of cooperation (see Fig. 1). We observe that *inclusive* approaches to interference negatively impact the mean frequency of cooperation if the individual endowments are not sufficient to turn defectors away from the temptation of defecting. By inclusive approaches, we imply high values for the threshold that determines the eligibility of investment (for POP and NEB schemes). If an external investor *hedges their bets*, targeting a wide spread of nodes (high threshold) with reduced individual endowments, they risk dooming cooperators. In such a scenario, we see the formation of cyclic patterns, ultimately allowing D players to exploit cooperators (see Fig. 2). In this way, an investor would be artificially allowing the survival of cooperators in clusters dominated by defectors, abetting the possibility of these sparsely connected clusters to take over larger formations which cannot easily be maintained by defectors. We note that some of these cyclic patterns eventually converge to a stable state, but the accumulated costs of interference at the end of these long-lasting patterns is prohibitively large.

In the presence of deterministic selection, this finding is mostly restricted to classical scale-free networks with low clustering (generated using the BA model), but relaxing the intensity of selection produces similar results even with more realistic levels of clustering (see Sect. 3.4). Social diversity changes the inherent nature of the problem of rewarding cooperators effectively. Previous results show the emergence of cooperation in heterogeneous networks [47,

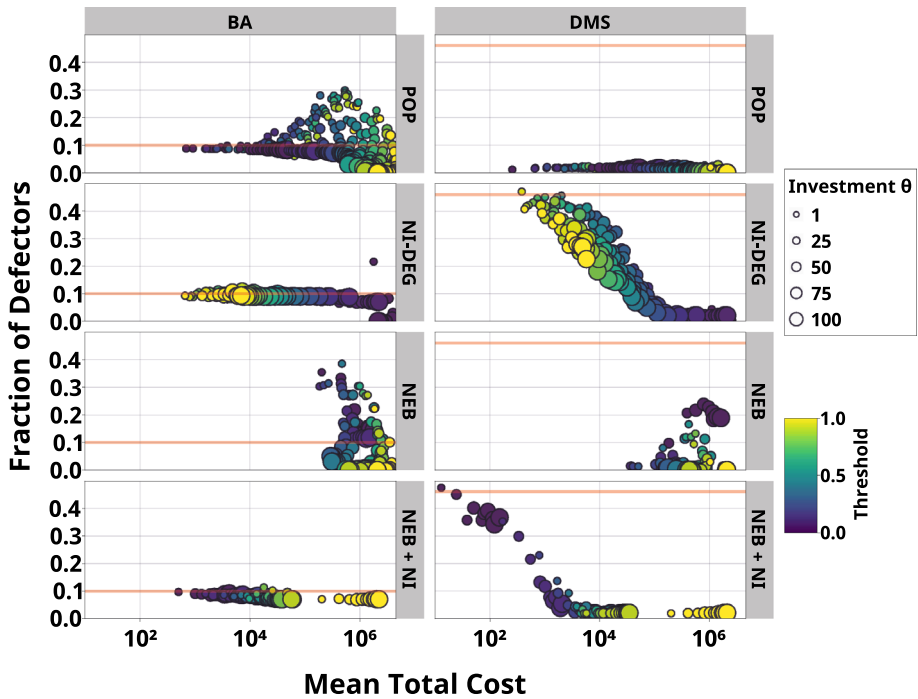


Fig. 1 Fraction of defectors as a function of the mean total cost for each scheme (deterministic update). The markers' size is determined by the individual investment θ (grouped to the nearest value), whereas the colour indicates the threshold. Points near the origin indicate the optimal solutions. The horizontal red lines indicate the baseline level of defection in the absence of rewards for either network type (i.e. BA or DMS). Parameters: $b = 1.8$; $N = 5000$

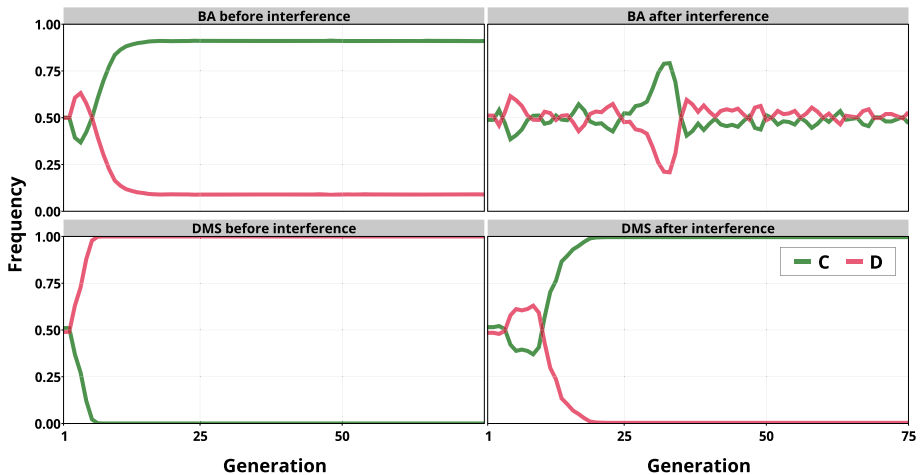


Fig. 2 Typical time evolution of cooperation, for $\theta = 5$, $p_C = 0.8$ (deterministic update). The left column shows the network without interference, while the right one shows the same network after population-based (POP) interference. Some configurations for BA resolve to full C, here we show the scenario in which they do not. Other parameters: $b = 1.8$; $N = 5000$

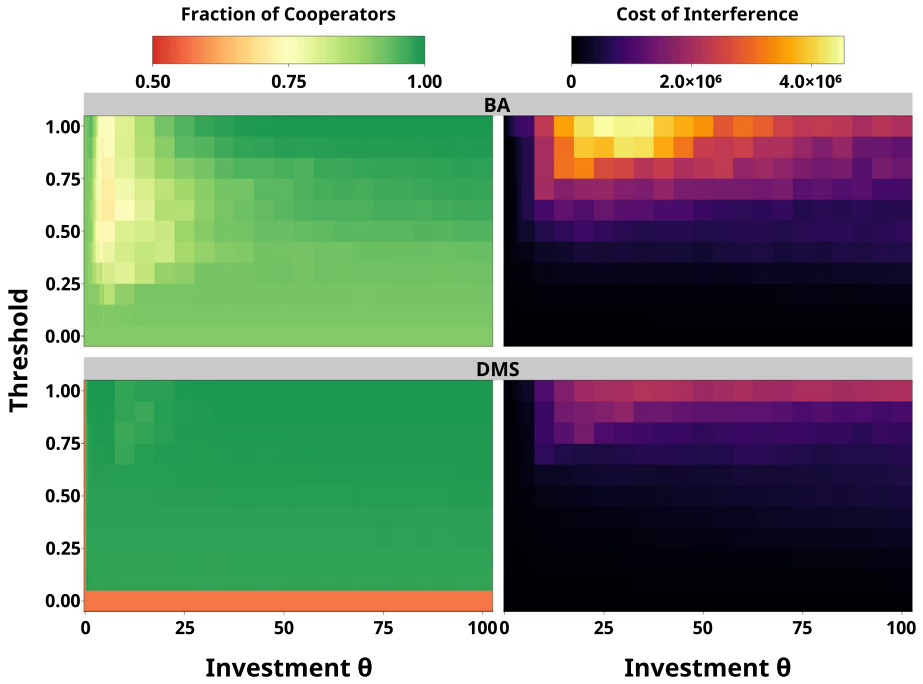


Fig. 3 Fraction of cooperation and total cost for population-based (POP) interference, using deterministic update. Parameters: $b = 1.8$; $N = 5000$

48] (shown also in horizontal red lines in Fig. 1). Compared to homogeneous (well-mixed) and structured populations, there is little improvement to be made in these settings. As the room for improvement narrows, the risk of acting to the detriment of cooperators increases. Individual benefactors prosper temporarily, but the recipients of their naivety are none other than the defectors who exploit them.

3.2 Clustering Reduces the Burden of Investment

Real-world networks have been observed to have higher levels of clustering than what normally occurs in typical scale-free networks [7, 54]. Nevertheless, several domains, such as the topology of the WWW, remain, in which the nodes are sparsely clustered [1, 4, 5]. Thus, it is important to design interference schemes which can target either type of scale-free networks, especially so if there exists a degree of uncertainty about the presence of clusters, or if measuring this factor is unfeasible. We have already mentioned the risks associated with inadequate reward mechanisms, but now we can turn to unveiling the benefits associated with social diversity and clustering in the quest towards engineering pro-social behaviour.

Highly clustered networks often have the most room to improve by receiving endowments (See Fig. 1). The initial distribution of players in the hubs of the network often determines whether the direction towards which the population will converge. Often, a small nudge can steer the population towards a desirable outcome (see Fig. 2). Moreover, this can easily be accomplished through a variety of disparate investment paradigms. For instance, metrics on the overall population (POP) can be used to guarantee maximal cooperation regardless

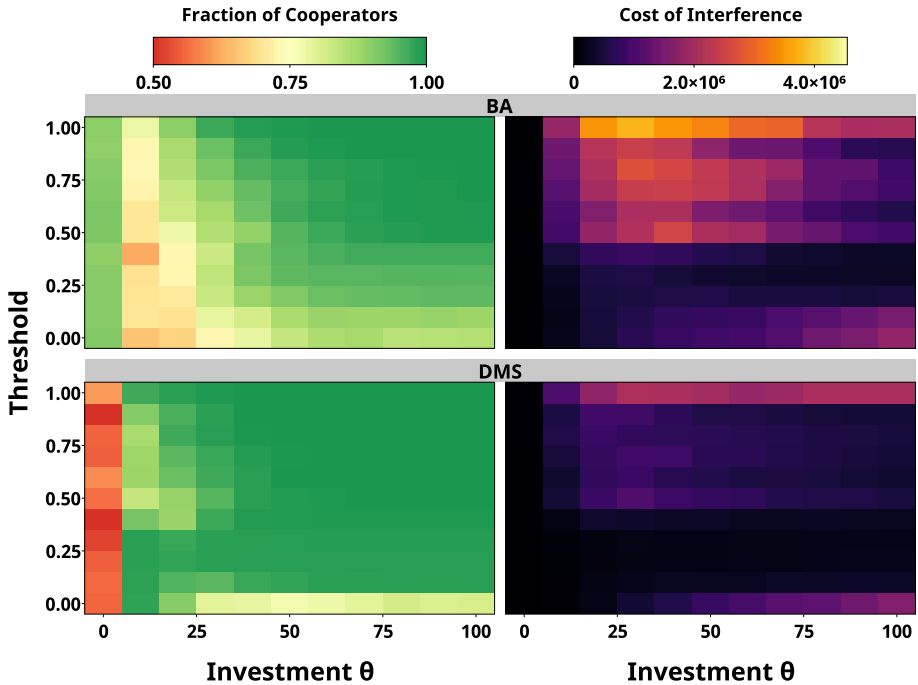


Fig. 4 Fraction of cooperation and total cost for local neighbourhood information (NEB) interference, using deterministic update. Parameters: $b = 1.8$; $N = 5000$

of how the endowments are distributed (See Fig. 3). With the reduction in the complexity of designing an effective scheme, we look towards cost and ways to reduce overspending. Overeager endowments can lead to total costs several orders of magnitude larger than those applied as a last resort. Indeed, even very small endowments applied to few surviving cooperators can *jumpstart* the formation of clusters resilient to invasion. Increasing the threshold for investment guarantees that more cooperators will be eligible for the endowments, thus exacerbating spending. Lowering this threshold guarantees that interference will only be triggered if desperately required. Investing in every cooperator as a last resort ensures pro-sociality.

Moreover, local observations can be used to ensure positive outcomes following a variety of pathways (see Fig. 4). In this case, an external decision-maker must target a range of intermediate values for the threshold. Previous results on structured populations showed that investing in cooperator neighbourhoods with exactly one defector was the optimal way of fostering cooperation [22]. In contrast, our findings suggest that the opposite is true for heterogeneous settings. Indeed, the least expensive routes towards cooperation are those with low or intermediate thresholds, suggesting that investors should focus their attention on ensuring only the survival of cooperators who are in danger of turning. For highly clustered networks, little investment is needed, and provided the threshold is not exceedingly low, maximal cooperation can be reached in any configuration, without unnecessary expenditure. Lowly clustered networks, on the other hand, require much more deliberate endowments to benefit from investment, with the added risk of causing cooperators to fall victim to exploitation as discussed previously in Fig. 4.

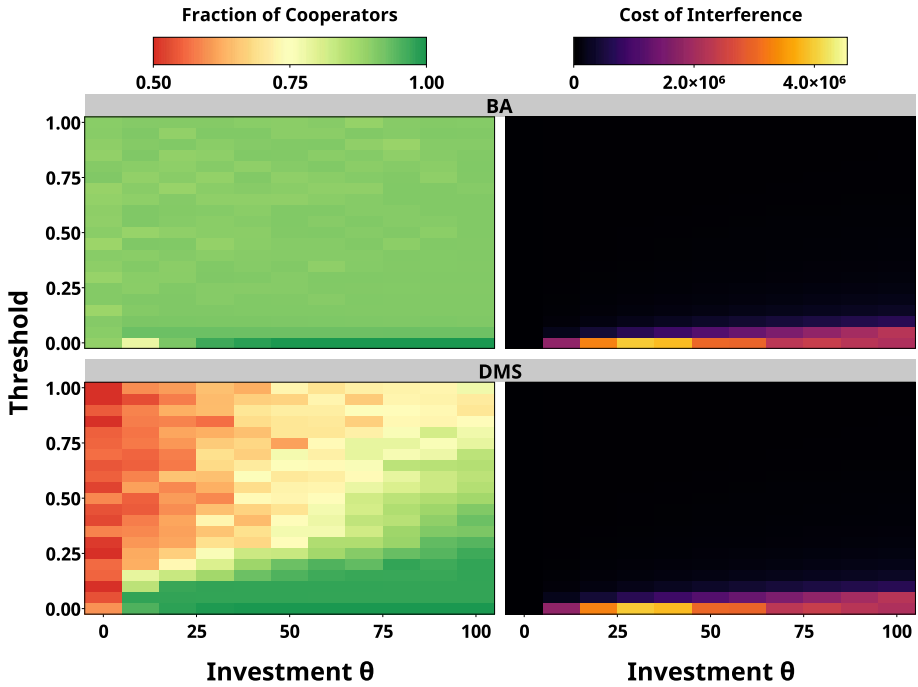


Fig. 5 Fraction of cooperation and total cost for node influence-based (NI) interference, using deterministic update. Parameters: $b = 1.8$; $N = 5000$

3.3 Heterogeneity and Network Characteristics Play a Key Role in the Design of Effective Investment Mechanisms

Assuming that information about a node’s influence can be easily gleaned by an external decision-maker, this can provide a partial solution to reducing the risk of deleterious interference. Although comparatively costly, this mechanism has the benefit of never succumbing to the exploitation of cooperators (see Fig. 5). Notwithstanding, the very nature of influential nodes in scale-free networks (i.e. power-law degree distribution) implies only exceedingly large endowments are sufficient to sway them. However, the number of cooperators who are eligible for investment is also small; on account of this, overall spending does not scale predictably with the endowment amount. We have previously mentioned that there exist some costs associated with information gathering, which we do not model or measure here. Hence, the assumption that information about influence is readily available suggests this method could prevail with respect to real-world budgeting.

We propose that combining several interference mechanisms can be an effective way of reducing spending while avoiding the pitfalls of pernicious investment. For instance, we might consider taking into account an agent’s influence as well as local observations. In Fig. 6, we explore this possibility, avoiding the least connected nodes (i.e. not investing in the bottom 5% of nodes with respect to degree centrality), and show that this reduces spending compared to either of the two interference schemes taken individually. These results suggest that hubs play an important role in the emergence of cooperation in highly clustered networks, but that they cannot be effectively used to improve outcomes in their lowly clustered counterparts.

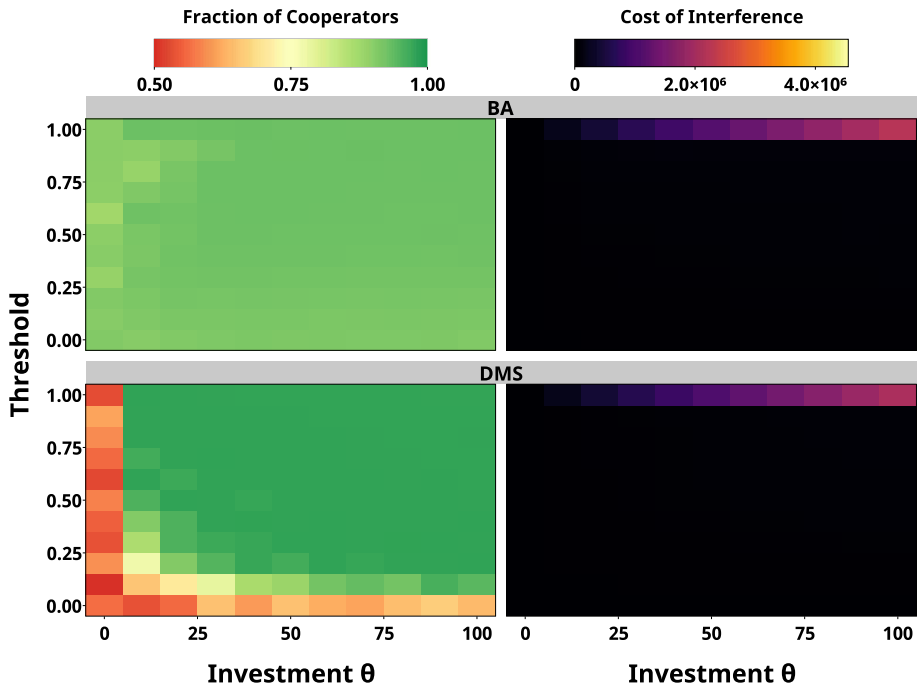


Fig. 6 Fraction of cooperation and total cost for a mixed interference scheme (NEB and NI), using deterministic update. We fix $c_I = 0.05$, avoiding investing into the least connected nodes (bottom 5%). Parameters: $b = 1.8$; $N = 5000$

Nevertheless, this integrated approach to interference eliminates the possibility of investment being detrimental to cooperation.

We note this conundrum between the two types of heterogeneous networks. Lowly clustered networks have little to benefit from investment, and much to lose if the external investor is negligent in their distribution of endowments. On the other hand, highly clustered networks have much to gain and little to lose, readily responding positively to any tactic, overspending being the only matter of discontent. As investment in the greater context of heterogeneous interactions is not trivial, it would therefore be prudent to first collect as much data on the nature of the network before deciding to distribute endowments. Uncertainty about social diversity or clustering carries the additional risk of selecting an improper policy of designing incentive schemes.

3.4 Stochastic Imitation Increases the Risk of Exploitation

Previously, we had shown that careless rewards might lead to an increase in defectors when interfering in BA networks under a deterministic update paradigm (see Sect. 3.1). Following a transition towards a more realistic, stochastic update rule [58], we observe a very similar phenomenon and moreover find that it is no longer limited to lowly clustered scale-free networks (see Fig. 7). Indeed, investing in DMS networks should be approached with the same due diligence as BA networks, and insufficient endowments often lead to the exploitation of cooperators. Relying solely on local information is most prone to damaging

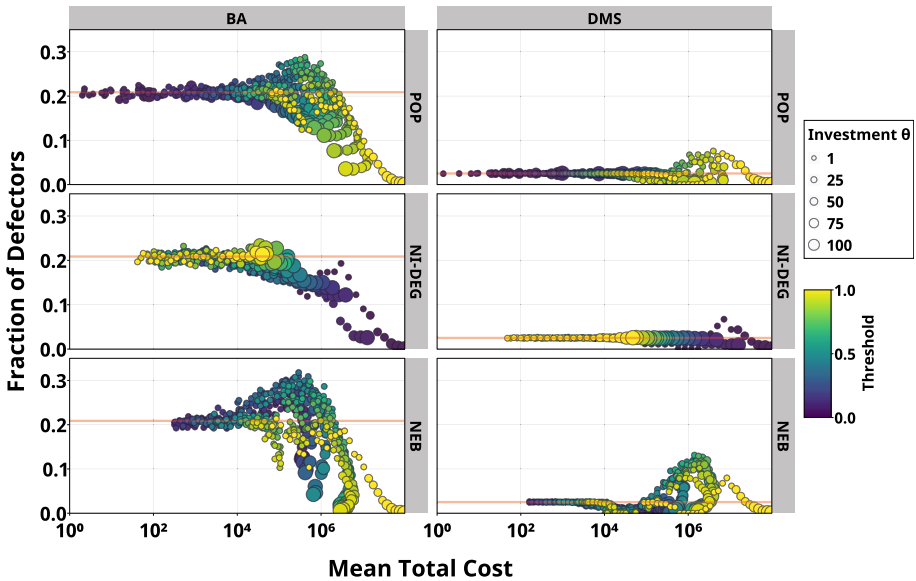


Fig. 7 Proportion of defectors as a function of the mean total cost for each scheme (stochastic update). The markers' size is determined by the individual investment θ (grouped to the nearest value), whereas the colour indicates the threshold. Points near the origin indicate the optimal solutions. The horizontal red lines indicate the baseline level of defection in the absence of rewards for either network type. Parameters: $b = 1.8$; $N = 2000$; $k = 0.1$

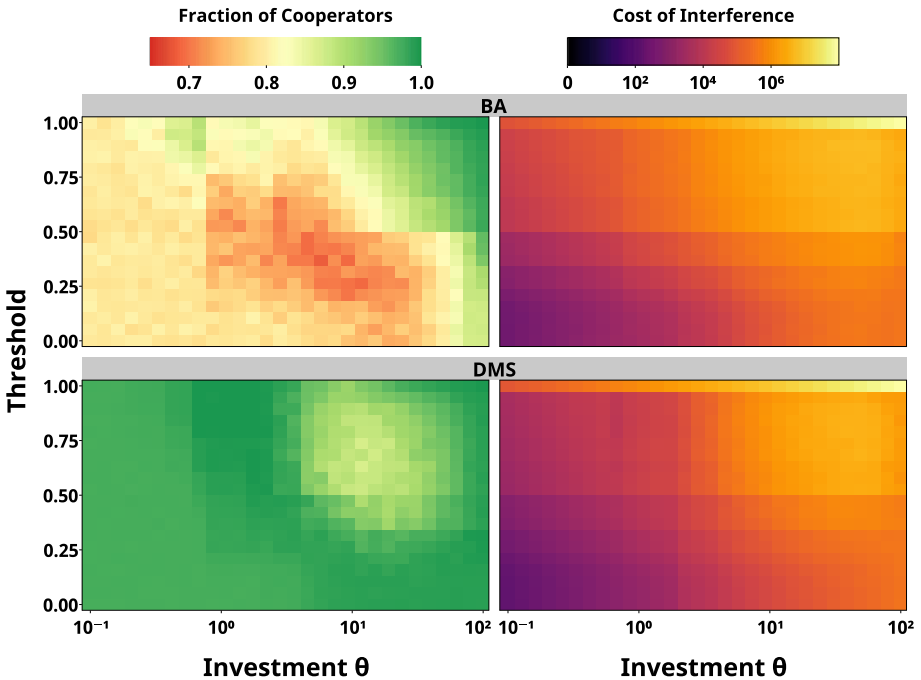


Fig. 8 Fraction of cooperation and total cost for local neighbourhood information (NEB) interference (stochastic update) using stochastic update. Parameters: $b = 1.8$; $N = 2000$; $k = 0.1$

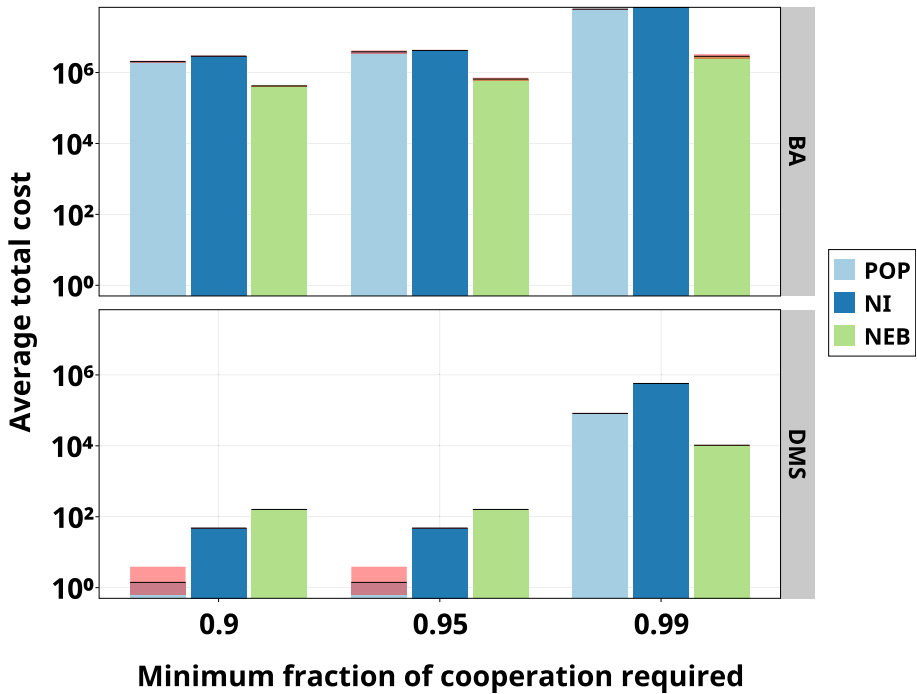


Fig. 9 Mean total costs for the most efficient combinations of threshold and investment amount (stochastic update) θ , using stochastic update. We intentionally avoid configurations in which no endowments are distributed, and select the configurations with the least possible cost for each minimal fraction of cooperation required. Error bars in light red show the standard deviation across all replicates of a configuration

cooperation, in spite of the level of complexity associated with this scheme and the amount of information required to enforce it (See Fig. 8).

Interestingly, stochastic imitation leads to a significant increase in baseline cooperation prior to interference in highly clustered networks and conversely a decrease in cooperation in classical scale-free networks (see horizontal lines in Fig. 7). Nevertheless, the findings discussed in the sections above remain robust. Although the potential gains to be had shift, causing BA networks to benefit from investment more than their highly clustered counterparts, the previous findings still apply in this setting. For instance, we find that DMS networks readily respond to investment, and are not as prone to the pitfalls which befall BA networks. In Fig. 9, we show that the most efficient interference schemes are consistently one (or more) order(s) of magnitude less costly at promoting cooperation in highly clustered networks, regardless of what potential gains the external decision-makers is aiming for.

3.5 Maximal Cooperation Gains Require Significant Endowments. Cost Efficiency is a Double-Edged Sword

Unlike homogeneous populations, heterogeneous interaction structures inherently provide a benefit to cooperators, something usually referred to as network reciprocity [47, 48]. In practice, this means investment does not lead to outcomes which differ significantly from the baseline. Furthermore, successful attempts at reaching a maximal level of cooperation (i.e. little to no defection) require a combination of large endowments and an investment scheme

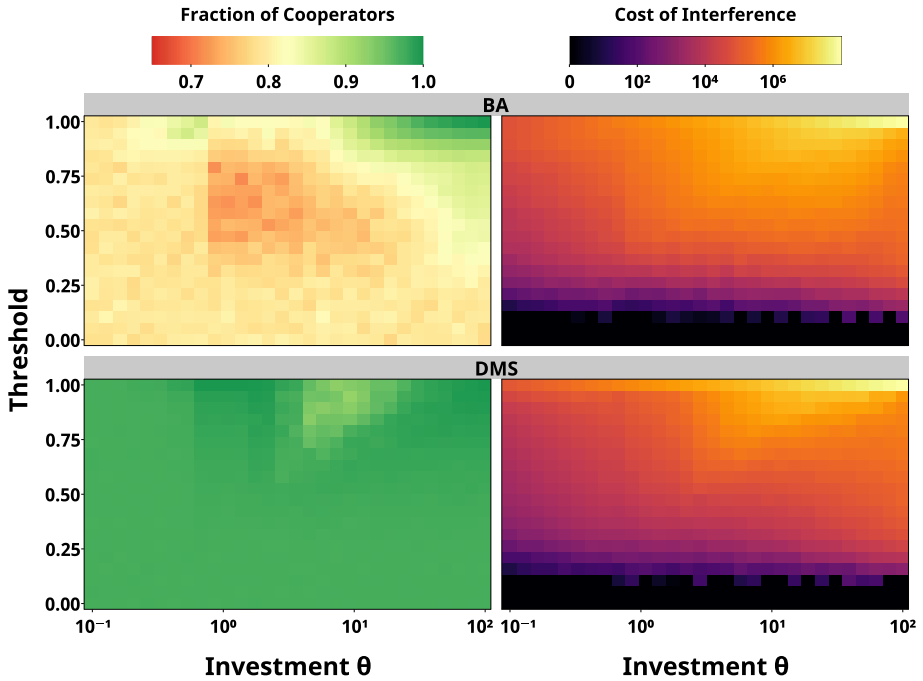


Fig. 10 Fraction of cooperation and total cost for population-based (POP) interference (stochastic update) using stochastic update. Parameters: $b = 1.8$; $N = 2000$; $k = 0.1$

which can target individuals at all levels of the network (see Fig. 7). Using population level metrics generally fails to improve outcomes unless virtually every cooperator is targeted (see Fig. 10). Equivalently, relying on degree centrality (i.e. how influential a node is) necessitates an egalitarian distribution of endowments, which naturally increases costs (See Fig. 11).

Local information (NEB), while risky, also has the potential to best improve outcomes while reducing costs, and this remains true for both BA and DMS networks, if maximal cooperation gains are required (See Fig. 9). Once again, we intuit the importance of acquiring detailed observations of information about the agents. This approach is a double-edged sword; it is simultaneously the optimal solution, as well as the most prone to errors in decision-making, leading the population to either the most perceived gains or the least (See Fig. 8). This seems to be another dilemma. Investing is risky, and it is likely for endowments to be ineffective or even produce negative results, but only significant sums of capital are likely to lead to desirable outcomes. Social diversity complicates this further, as there exists a great degree of inequality between individuals, and potential errors in decision-making make investment precarious.

4 Discussion

In summary, we have studied how optimally an external decision-maker could incentivise a population of autonomous agents facing a cooperative dilemma to fulfil a coveted collective state. We build on a previous account which identified the most effective mechanisms to foster cooperative scenarios in spatially distributed systems in regular graph structured

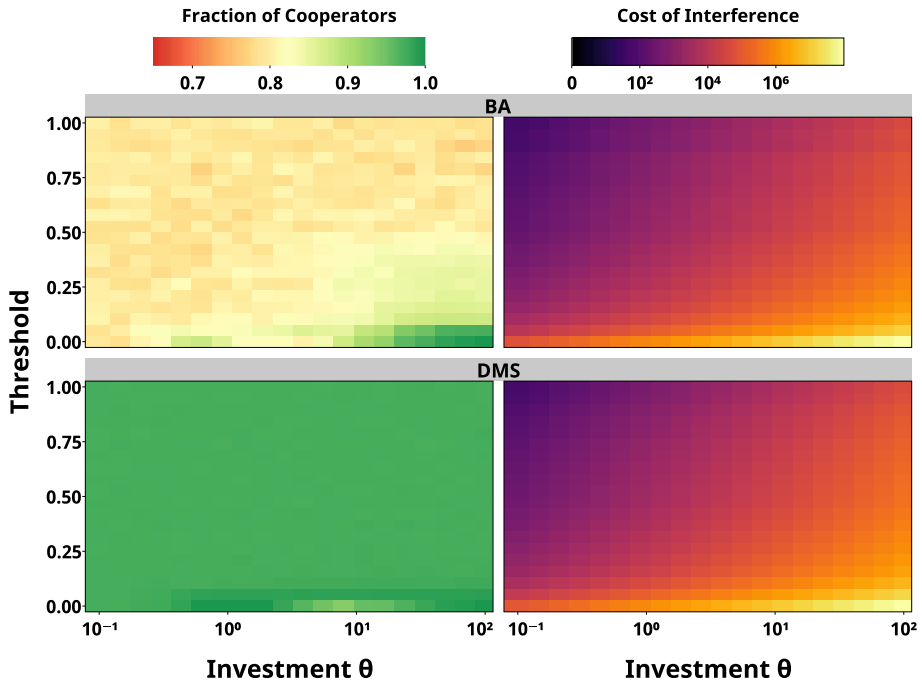


Fig. 11 Fraction of cooperation and total cost for node influence-based (NI) interference (stochastic update) using stochastic update. Parameters: $b = 1.8$; $N = 2000$; $k = 0.1$

populations of agents, but instead we consider two popular models of scale-free networks of contacts. In particular, we examined if the insights set out in the context of regular graphs remain applicable to heterogeneous models, as well as explore an additional avenue of interference enabled by the variance in node connectivity. To address these issues, we combined an evolutionary game theoretic model with several incentive mechanisms in two types of pre-generated networks characterised by preferential attachment, with different clustering coefficients. We argue that this problem cannot be solved trivially and we show that transitivity (i.e. the global clustering coefficient) should be the driving force behind the choice of an interference mechanism in promoting cooperation in heterogeneous network structures, as well as its application.

In this work, we introduce several incentive mechanisms which are defined formally and mathematically. We note that they do not have to be defined as such and, in fact, have many real-life counterparts which are often employed by institutions and investors. For instance, POP-based metrics describe cooperation observed at a global scale. If we consider the Great Recession, or the recent COVID-19 pandemic, an institution might only need to look at the overall state of the economy, or the spread of an infection, before deciding that action is required. Neighbourhood-based metrics represent local schemes, which are almost ubiquitous when considering social inequality. Whether it is housing schemes, incentives to stop smoking, homelessness, education, etc., local governments often decide to invest based on the level of economic and social deprivation in a specific area, and that is precisely what we have tried to capture with NEB-based schemes. Finally, we have looked at centrality (influence) metrics. If we consider social media, a company might wish to use influencers to market its products, or an institution might decide to specifically target someone in the public

eye in order to increase the visibility of its incentives, whether they were positive or negative. The mechanisms we have chosen are by no means exhaustive choices, but they serve as a fundamental starting point to our discussion, and they are arguably the most common and most easily implementable mechanisms that we observe in the real world.

We find that impetuously rewarding cooperators can lead to cyclic patterns which damage cooperation in the long run, enabling the exploitation of cooperators to the benefit of defectors. We argue that detailed information gathering about the networks and agents prior to the distribution of endowments can prevent these mistakes. Using two social learning paradigms, we show the robustness of these findings and observe that clustering lowers the risk of deleterious investment, easing the strictness of distributing incentives. Moreover, we show that ignoring lowly connected individuals leads to unprofitable and even futile intervention irrespective of network transitivity.

Our comparison between the two types of scale-free networks provides valuable insights regarding the importance of clustering in the outcome of cooperation. We find that a large clustering coefficient allows for successful, cost-effective interference, indeed even when partly disregarding a full comprehension of the population and its tendencies. Furthermore, transitivity lessens the burden on external investors, lowering the total cost required to enforce cooperation. These results are of particular interest, given that most SF networks portray high clustering, such as in the case of social ties where friends are likely to be friends of each other [33]. This scope encompasses heterogeneous scenarios inhibited by spatial constraints (e.g. in highly urbanised areas or even the allotment of rangelands such as pastures), where high clustering is also imposed.

Transitioning towards a more realistic, stochastic imitation rule [58], we measure a shift between the two network types, whereby lowly clustered networks prescribe a greater need for investment, and vice versa. Maximal cooperation gains in either paradigm can generally be achieved using large individual endowments. Notwithstanding, highly clustered networks respond more readily to interference, and we provide several insights about ways in which cost could be reduced further.

An important question in models of institutional incentives is that of setting up and maintaining the incentive budget. Considering who should contribute to the incentive budget is a social dilemma in itself, and addressing this second-order social dilemma has been identified as a challenging research problem. Several solutions have been identified, including pool incentives with second-order punishments [39, 53], democratic decisions [27], commitment formation [21, 25, 49] and hybrid incentives [10, 20]. This work does not aim to address this issue, focusing instead on how to optimise the spending from a given budget by exploiting network properties and information gathering. However, it would be interesting to study the co-evolutionary institutional formation with different interference strategies and individual strategic behaviours. Moreover, in this work, we do not consider the possibility of detecting the existence of a certain type of interference from an external party. In reality, individuals could be aware of active interference and react by changing their behaviour, either to become suitable candidates for reward or to avoid sanctions [12].

In summary, we have shown that it is crucial to investigate both the underlying network and behavioural trends before distributing endowments. Previously, these findings were restricted to very particular combinations of parameters, but stochastic dynamics highlight the pitfalls of inconsiderate incentives in a variety of settings. As a result of such careless interference, defector communities grow before eventually collapsing due to their inability to reciprocate. We have identified several key metrics which can be used to mitigate these risks and ensure

positive social outcomes. Our findings highlight the complexity of designing effective and cogent investment policies in socially diverse populations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Albert R, Jeong H, Barabási A-L (1999) Diameter of the World-Wide Web. *Nature* 401(6749):130–131
2. Allen B, Lippner G, Chen YT, Fotouhi B, Momeni N, Yau ST, Nowak MA (2017) Evolutionary dynamics on any population structure. *Nature* 544(7649):227–230
3. Andras P, Esterle L, Guckert M, Han TA, Lewis PR, Milanovic K, Payne T, Perret C, Pitt J, Powers ST, Urquhart N, Wells S (2018) Trusting intelligent machines: deepening trust within socio-technical systems. *IEEE Technol Soc Mag* 37(4):76–83
4. Barabasi A-L (2014) *Linked-how everything is connected to everything else and what it means*. F. Perseus Books Group
5. Barabási A-L (2016) *Network science*. Cambridge University Press, Cambridge
6. Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512
7. Barrat A, Pastor-Satorras R (2005) Rate equation approach for correlations in growing network models. *Phys Rev E* 71(3):36127
8. Bloch F, Jackson MO, Tebaldi P (2019) Centrality measures in networks. Available at SSRN 2749124
9. Boldi P, Vigna S (2014) Axioms for centrality. *Internet Math* 10(3–4):222–262
10. Chen X, Sasaki T, Brännström Å, Dieckmann U (2015) First carrot, then stick: how the adaptive hybridization of incentives promotes cooperation. *J R Soc Interface* 12(102):20140935
11. Cimpanu T, Di Stefano A, Perret C, Han TA (2023) Social diversity reduces the complexity and cost of fostering fairness. *Chaos Solitons Fractals* 167:113051
12. Cimpanu T, Han TA (2020) Making an example: signalling threat in the evolution of cooperation. In: 2020 IEEE Congress on Evolutionary Computation (CEC), pp 1–8. IEEE
13. Cimpanu T, Perret C, Han TA (2021) Cost-efficient interventions for promoting fairness in the ultimatum game. *Knowl-Based Syst* 233:107545
14. Dafoe A, Bachrach Y, Hadfield G, Horvitz E, Larson K, Graepel T et al (2021) Cooperative ai: machines must learn to find common ground. *Nature* 593(7857):33–36
15. Dall'Asta L, Baronchelli A, Barrat A, Loreto V (2006) Nonequilibrium dynamics of language games on complex networks. *Phys Rev E* 74(3):36105
16. Dorogovtsev S (2010) *Complex networks*. Oxford University Press, Oxford
17. Dorogovtsev SN, Mendes JFF, Samukhin AN (2000) Structure of growing networks with preferential linking. *Phys Rev Lett* 85(21):4633
18. Duong, M, Durbac C, Han T (2022) Cost optimisation of hybrid institutional incentives for promoting cooperation in finite populations. arXiv preprint [arXiv:2212.08823](https://arxiv.org/abs/2212.08823)
19. Duong MH, Han TA (2021) Cost efficiency of institutional incentives in finite populations. In: Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences
20. Góis AR, Santos FP, Pacheco JM, Santos FC (2019) Reward and punishment in climate change dilemmas. *Sci Rep* 9(1):1–9
21. Han TA (2022) Institutional incentives for the evolution of committed cooperation: ensuring participation is as important as enhancing compliance. *J R Soc Interface* 19(188):20220036
22. Han TA, Lynch S, Tran-Thanh L, Santos FC (2018) Fostering cooperation in structured populations through local and global interference strategies. In: *IJCAI-ECAI'2018*, pp 289–295. AAAI Press
23. Han TA, Pereira LM, Lenaerts T (2019) Modelling and influencing the AI bidding war: a research agenda. In: Proceedings of the AAAI/ACM conference AI, Ethics and Society, p 5–11
24. Han TA, Perret C, Powers ST (2021) When to (or not to) trust intelligent machines: insights from an evolutionary game theory analysis of trust in repeated games. *Cogn Syst Res* 68:111–124

25. Han TA, Santos FC, Lenaerts T, Pereira LM (2015) Synergy between intention recognition and commitments in cooperation dilemmas. *Sci Rep* 5(9312):1–7
26. Han TA, Tran-Thanh L (2018) Cost-effective external interference for promoting the evolution of cooperation. *Sci Rep* 8:15997
27. Hilbe C, Traulsen A, Röhl T, Milinski M (2014) Democratic decisions establish stable authorities that overcome the paradox of second-order punishment. *PNAS* 111(2):752–756
28. Hofbauer J, Sigmund K (1998) *Evolutionary games and population dynamics*. Cambridge University Press, Cambridge
29. Levin SA (2000) Multiple scales and the maintenance of biodiversity. *Ecosystems* 3(6):498–506
30. Maynard Smith J (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge
31. McAvoy A, Allen B (2021) Fixation probabilities in evolutionary dynamics under weak selection. *J Math Biol* 82(3):14
32. McAvoy A, Rao A, Hauert C (2021) Intriguing effects of selection intensity on the evolution of prosocial behaviors. *PLoS Comput Biol* 17(11):e1009611
33. Newman M (2018) *Networks*, 2nd edn. Oxford University Press, Oxford
34. Newman MEJ (2003) The structure and function of complex networks. *SIAM Rev* 45(2):167–256
35. Newman MEJ (2008) The mathematics of networks. *New Palgrave Encycl Econ* 2(2008):1–12
36. Nowak MA, May RM (1992) Evolutionary games and spatial chaos. *Nature* 359(6398):826–829
37. Paiva A, Santos FPFC, Santos FPFC (2018) Engineering pro-sociality with autonomous agents. *Proc AAAI Conf Artif Intell* 32:7994–7999
38. Penn AS, Watson RA, Kraaijeveld A, Webb J (2010) Systems Aikido—a novel approach to managing natural systems. In: *Proceedings of the ALIFE XII Conference*, pp 577–580. MIT press
39. Perc M (2012) Sustainable institutionalized punishment requires elimination of second-order free-riders. *Sci Rep* 2(1):1–6
40. Perc M, Jordan JJ, Rand DG, Wang Z, Boccaletti S, Szolnoki A (2017) Statistical physics of human cooperation. *Phys Rep* 687:1–51
41. Perc M, Szolnoki A (2010) Coevolutionary games—a mini review. *Biosystems* 99(2):109–125
42. Poncela J, Gómez-Gardenes J, Floría LM, Moreno Y (2007) Robustness of cooperation in the evolutionary prisoner’s dilemma on complex networks. *New J Phys* 9(6):184
43. Raghunandan MA, Subramanian CA (2012) Sustaining cooperation on networks: an analytical study based on evolutionary game theory. In: *AAMAS’12*, volume 12, pp 913–920. Citeseer
44. Rand DG, Nowak MA (2013) Human cooperation. *Trends Cogn Sci* 17(8):413–425
45. Rand DG, Nowak MA, Fowler JH, Christakis NA (2014) Static network structure can stabilize human cooperation. *Proc Natl Acad Sci USA* 111(48):17093–17098
46. Santos FC, Pacheco JM (2005) Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys Rev Lett* 95:98104
47. Santos FC, Pacheco JM, Lenaerts T (2006) Evolutionary Dynamics of Social Dilemmas in Structured Heterogeneous Populations. *Proc Natl Acad Sci USA* 103:3490–3494
48. Santos FC, Santos MD, Pacheco JM (2008) Social diversity promotes the emergence of cooperation in public goods games. *Nature* 454:214–216
49. Sasaki T, Okada I, Uchida S, Chen X (2015) Commitment to cooperation and peer punishment: its evolution. *Games* 6(4):574–587
50. Saxena A, Iyengar S (2020) Centrality measures in complex networks: a survey. *arXiv preprint arXiv:2011.07190*
51. Sigmund K (2010) *The calculus of selfishness*. Princeton University Press, Princeton
52. Sigmund K, Hauert C, Nowak M (2001) Reward and punishment. *Proc Natl Acad Sci* 98(19):10757–10762
53. Sigmund K, Silva HD, Traulsen A, Hauert C (2010) Social learning promotes institutions for governing the commons. *Nature* 466:7308
54. Su J, Sharma A, Goel S (2016) The effect of recommendations on network structure. *25th International World Wide Web Conference. WWW 2016*:1157–1167
55. Su Q, McAvoy A, Plotkin JB (2022) Evolution of cooperation with contextualized behavior. *Sci Adv* 8(6):eabm6066
56. Szabó G, Fáth G (2007) Evolutionary games on graphs. *Phys Rep* 446(4–6):97–216
57. Szabó G, Toke C (1998) Evolutionary Prisoner’s Dilemma game on a square lattice. *Phys Rev E* 58:69–73
58. Traulsen A, Nowak MA, Pacheco JM (2006) Stochastic dynamics of invasion and fixation. *Phys Rev E* 74:11909
59. Wang S, Chen X, Szolnoki A (2019) Exploring optimal institutional incentives for public cooperation. *Commun Nonlinear Sci Numer Simul* 79:104914

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