

Community connectivity and heterogeneity: clues and insights on cooperation on social networks.

Sergi Lozano · Alex Arenas ·
Angel Sánchez

Received: date / Accepted: date

Abstract While studies on the emergence of cooperation on structured populations abound, only few of them have considered real social networks as the substrate on which individuals interact. As has been shown recently [Lozano *et al.*, PLoS ONE **3**(4): e1892], understanding cooperative behavior on social networks requires knowledge not only of their global (macroscopic) characteristic, but also a deep insight on their community (mesoscopic) structure. In this paper, we look at this problem from the viewpoint of the resilience of cooperation, in particular when there are directed attacks (insertion of pure defectors) at key locations in the network. We present results of agent-based simulations showing strong evidence that the resilience of social networks is crucially dependent on their community structure, ranging from no resilience to robust cooperative behavior. Our results have important implications for

This work was supported by Ministerio de Educación y Ciencia (Spain) under grants FIS2006-13321, MOSAICO and NAN2004-9087-C03-03 and by Comunidad de Madrid (Spain) under grants UC3M-FI-05-007 and SIMUMAT-CM. S.L. was supported by URV through a FPU grant and by the EU Integrated Project IRRIS (027568).

Sergi Lozano
ETH Zurich, Swiss Federal Institute of Technology, UNO D11, Universitätstr. 41, 8092 Zurich, Switzerland
E-mail: slozano@ethz.ch

Alex Arenas
Universitat Rovira i Virgili, Avda. Països Catalans 26, 43007 Tarragona, Spain
Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Zaragoza, Spain
E-mail: alexandre.arenas@urv.cat

Anxo Sánchez
Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), Universidad de Zaragoza, Zaragoza, Spain
Grupo Interdisciplinar de Sistemas Complejos (GISC), Departamento de Matemáticas, Universidad Carlos III de Madrid, Leganés, Spain
E-mail: anx@math.uc3m.es

the understanding of how organizations work and can be used as a guide for organization design.

Keywords Cooperation · Social networks · Community structure · Resilience · Prisoner's Dilemma

PACS 87.23.Ge and 89.65.Ef · 89.65Gh · 64.60.Aq

1 Introduction

How cooperation emerges in human groups is one of the most intriguing puzzles of modern science. Indeed, in a world of rational individuals, cooperation, understood as helping others at a personal cost, is hard to reconcile with the *homo economicus* paradigm, guided only by the goal of maximizing her utility function. Notwithstanding, the very fact that complex human societies exist, with division of labor and welfare programs shows that cooperation does take place, hence the exciting challenge of its explanation.

The roots of cooperative behavior can be traced back to a number of different reasons, ranging from kin selection to group selection through several reciprocity mechanisms [1]. In this paper, we will focus on network reciprocity, a mechanism first suggested by Nowak and May [2]. In this scenario, the individuals of a population occupy the vertices of a graph, that in the specific context of human populations represents the social network of contacts (friendship, work, family, etc.). The edges determine who interacts with whom. In this setting, cooperators can prevail by forming network clusters, where they help each other. The term network reciprocity generalizes Nowak and May's first proposal of representing the social network by a simple model (a square lattice with next-neighbor and next-nearest-neighbor interaction) to account for the fact that actual social relationships may be better described by complex networks [3].

Following the original proposal, a great deal of research has been undertaken on evolutionary game theory on graphs [4]. However, none of these studies deals with true social networks, as they are all based on different types of artificial models reproducing only some of real social networks characteristics. To our knowledge, the main contribution to the study of the emergence of cooperation in social networks is a recent work of ours [5], where we have considered static empirical social networks as a support for the local interactions in the framework of imitation models. The main conclusion of that work is that cooperation on real social networks is a complex issue depending on the combination of the effects of several structural features and, consequently, that any approximation to the evolution of cooperation in social networks based on the generalization of only one of these structural features is far too simplistic. The features emerging from our research as the key to understand cooperative behavior on social networks are the communities, or sub-graphs consisting of nodes densely interconnected, which can be considered as a kind

of mesoscopic structure bridging the properties of individual agents with the global characteristics of the network (e.g., the degree distribution).

In this paper, we go beyond that previous work to consider the resilience of the observed cooperative behavior. Being formed by human beings, social networks are certainly vulnerable to external factors that may lead to non-cooperative behavior of some individuals. In turn, the influence of these individuals may have as a consequence the loss of the cooperative functioning of the whole network. Therefore, even if in our previous study we established that cooperators may outperform defectors and give rise to a very large level of cooperation in an evolutionary setting, it is important to analyze the resilience of such emergent behavior. As we will see below, our results show once again that the community structure governs the response of the social network to external conditioning factors, and that the resilience is very different depending on the network under consideration. To better present our results, we begin by summarizing the results in [5], introducing alongside the theoretical framework for our study, namely the most strict social dilemma, Prisoner's Dilemma. Building on those results, we study the resilience properties by means of an evolutionary agent based model, and present our results in Sec. 3. In order to understand our computer simulation results, Sec. 4 presents a simple, toy-model calculation that allows to explain why resilient behavior may or may not arise in terms of the community structure of the network. Finally, the last section summarizes the work and discusses its implications.

2 Model and previous results

As already mentioned above, recent work of ours [5] has uncovered the key role of the community structure on cooperation on social networks. Our research program took us from the analysis of cooperative behavior by means of agent-based simulations of the Prisoner's Dilemma on two real social networks to the modelling and application of our conclusions through the identification of the most relevant features of the networks: community interconnectivity and their internal heterogeneity. This section contains a brief summary of the results reported in [5] as relevant background for our research on the resilience of cooperation, which constitutes the main body of the paper.

2.1 Model

The goal of our research was to study cooperation on social networks and how it differs from the simulations on artificial models, such as scale-free or small-world networks. To that end, we used two social datasets obtained by sampling two kinds of interpersonal relational activities. The first one was extracted from the email traffic between members of University Rovira i Virgili (Tarragona, Spain; *email network* from now on): Nodes represent individual email addresses, and are connected by (undirected) when bidirectional communication among them is detected (at least one email in each direction within

Table 1 Statistical properties of e-mail and PGP networks. N is the number of nodes of the giant component of the network considering only those links that are bidirectional (indicating mutual acquaintance between nodes). $\langle k \rangle$ and $\langle C \rangle$ are the average degree and clustering coefficient, correspondingly. $P(k)$ is the degree distribution (best fit to the data using least squares method). $C(k)$ corresponds to the degree-dependent clustering coefficient, that is, the result of averaging C over vertices of degree k . Finally, r stands for the assortativity coefficient.

| Network | ref. | N | $\langle k \rangle$ | $P(k)$ | $\langle C \rangle$ | $C(k)$ | r |
|---------|------|-------|---------------------|--|---------------------|---------------------------|-------|
| email | [6] | 1133 | 9.62 | $\sim \exp^{-k/9.2}$ | 0.25 | $\sim 0.44 - 0.08 \ln(k)$ | 0.078 |
| PGP | [7] | 10680 | 4.55 | $\sim \begin{cases} k^{-2.63} & \text{if } k < 40 \\ k^{-4.0} & \text{if } k > 40 \end{cases}$ | 0.26 | $\sim 0.67 - 0.08 \ln(k)$ | 0.238 |

a short time period) [6]. Our second social network consists of the users of the "Pretty-Good-Privacy" encryption algorithm (*PGP network*, from now on), where links trace trust relationships between those persons who sign each other's public keys [7]. The networks were chosen because they arguably represent social relationships between individuals, something that may be more difficult to assess in other sets, such as co-authorship of scientific papers [3] where it can even be the case that two co-authors have never spoken to each other. On the other hand, their statistical properties (see table 1 for the data) were similar, in the sense that the two networks exhibited a typical degree (i.e., they were degree-homogeneous as opposed to scale-free) and similar global parameters, such as clustering, an important feature of social networks.

In the context of these two networks, we model the social dilemma of cooperation at a personal cost by making the individuals play a Prisoner's Dilemma (PD) [8]. The PD embodies a stringent form of social dilemma, namely a situation in which individuals can benefit from mutual cooperation but they can do even better by exploiting cooperation of others. To be specific, the two players in the PD can adopt either one of two strategies: cooperate (C) or defect (D). Mutual cooperation results in a benefit R for both players, whereas mutual defection yields them P . When a cooperator meets a defector, the latter receives T , whereas the former gets S . The PD is completely specified by imposing $T > R > P > S$. In either case, it is better for both players to play D, in spite of the fact that mutual cooperation would yield higher benefits for them, hence the dilemma. The PD is a particularly appropriate choice because it applies to very many different situations of interactions between human individuals or collectives [9,10] and, interestingly, even to microorganisms [11]. For our work, we chose the same parameterization of the PD as in [2]: $T = b > 1, R = 1, S = P = 0$.

Our agent-based simulations proceed as follows: We place individuals on the nodes of our social networks. At every time step, individuals play the game with their nearest neighbors. After receiving their payoffs according to the PD, all individuals update their strategy synchronously for the next round, by imitate-the-best (also called unconditional imitation) dynamics [2]: They look in their neighborhood for players whose payoff is higher than their own and, if

there is any, they adopt the strategy that led to that highest payoff (randomly chosen in case of a tie). We then repeat the process and let the simulation run until the density of cooperators in the lattice reaches an asymptotic average value or else it becomes 0 or 1 (note that these two states, corresponding to full defection and full cooperation, are absorbing states of the dynamics because there are not mutations).

2.2 Previous results: the influence of communities

In [5], we carried out a thorough simulation program of the above described model on our two social networks. Our results can be summarized in two main remarks:

Remark 1. The behaviour of the two networks was completely different. The cooperation level on the email network was a decreasing function of b , going from values very close to unanimous cooperation for $b \gtrsim 1$ to about a 15% for b close to 2. On the contrary, the PGP network presented an almost constant cooperation level, with a variation of a 10% at most in all the range of b values, except for $b = 2$.

Remark 2. Neither of our two networks seemed to fit in any of the categories previously reported in the literature for the behavior of the PD (see [4] and references therein).

These two remarks, along with the fact that from the macroscopic (global, statistical) viewpoint both networks are very similar, led us to conclude that those global properties were not determinant for the opposite behaviors observed. On the other hand, at the individual level the results could not be explained as well, and therefore we were forced to look at a mesoscopic level, i.e., to their community structure. It is difficult to give a precise definition of community, but for the present purpose we can understand "communities" as subsets of nodes of the network which are very densely connected among themselves while their connections to the rest of the network are more sparse. More rigorous definitions are of course available (see [3] and references therein). For our work, we resorted to a community-finding algorithm based on a divisive procedure proposed by one of the authors, which relies on Extremal Optimization (EO) heuristics [12]. In Fig. 1, we present graphically the results obtained to facilitate the comparison among the two cases. Careful analysis of the results led us to identify two features that characterize the community structure and allow to discriminate between our two networks:

Inter-community connectivity (IC). The results corresponding to this structural feature are presented at the top of Fig. 1. Each node corresponds to a community, and a link between two nodes denotes cross-relations. In addition,




Fig. 1 Community structures of the email and PGP networks. Top: Community structures of the email (left) and PGP (right) networks. Nodes correspond to communities (where size is proportional to their number of members) and links represent cross-connections (where width corresponds to the number of inter-connections). Bottom: Typical examples of the communities detected in the email (left) and PGP (right) networks. Solid links join nodes of the community, dashed links join this community with others. Reproduced from [5] with permission.

the size of nodes and links gives information about community size and number of cross-links, respectively. It is evident from the plot that communities in the email network are densely interconnected, and sparsely interconnected in the PGP network.

Intra-community heterogeneity (IH). The internal structure of communities in both networks also presents important differences. In Fig. 1 (bottom), we plot the aspect of representative communities of the email and PGP networks, respectively. From the plot, the differences in the internal structure are clear: the email communities present a very homogeneous structure when compared with the heterogeneity of PGP communities. In the following, we will call *local hubs* the nodes in the PGP networks responsible for the very high internal heterogeneity.

Taking into account the large differences both at the intra-community heterogeneity (IH) and the inter-community connectivity (IC) between the e-mail and the PGP networks, we formulated an hypothesis: The robustness of cooperation observed in the PGP network is due to its low IC and high IH, whereas the fact that cooperation only arises for low b in the email network is caused

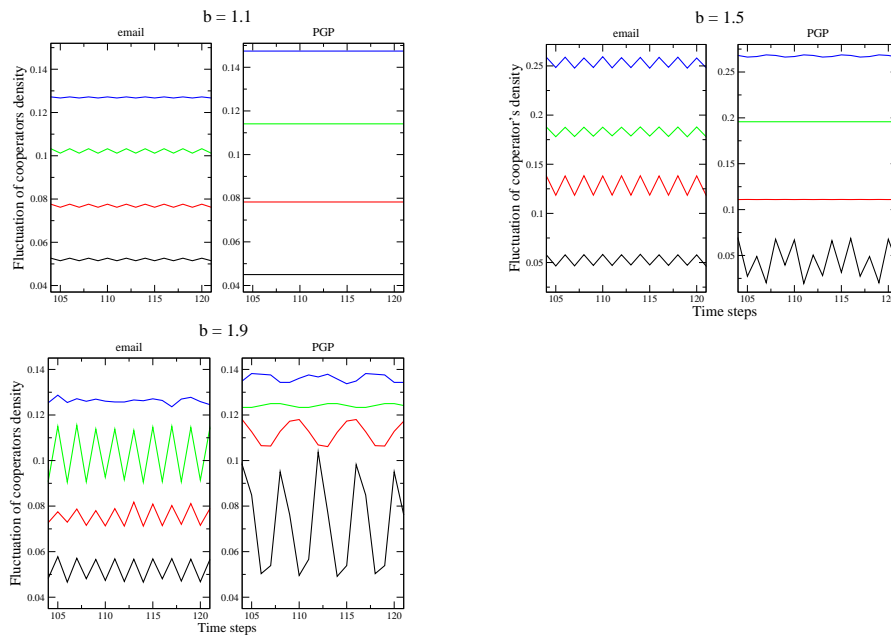


Fig. 2 Temporal evolution of cooperation within email and PGP communities. Fluctuation of cooperators' density within four sample communities of both empirical networks for three different values of b . In general, we observe that, in general, communities in email network present some kind of synchronization, indicating a high degree of interdependence (high IC). On the contrary, PGP communities show regular behaviors with different periods, revealing an strong influence of internal structural properties of each community (such as their high IH).

by its high IC and low IH.

Significantly, this conjecture fits nicely with certain traits of cooperation's temporal evolution within each community of both empirical networks, specially in terms of stability and coordination of behaviours across communities (see Fig 2 for some examples). Furthermore, a proper confirmation can be found in [5] where we studied four statistically significant synthetic test networks with the same number of nodes and the same number of communities, corresponding to four extreme configurations obtained by the combination of low and high values of the IH and IC. Two of the networks presented the configuration corresponding to the email network (low IH and high IC) and the PGP one (high IH and low IC). The other two cases, low IH and low IC, and high IH and high IC, corresponded to intermediate configurations between the former ones. Simulations on these four synthetic cases reproduced the opposite behaviours observed for the email and PGP networks, thus supporting our conclusion of the relevance of the community structure (see Fig. 3).

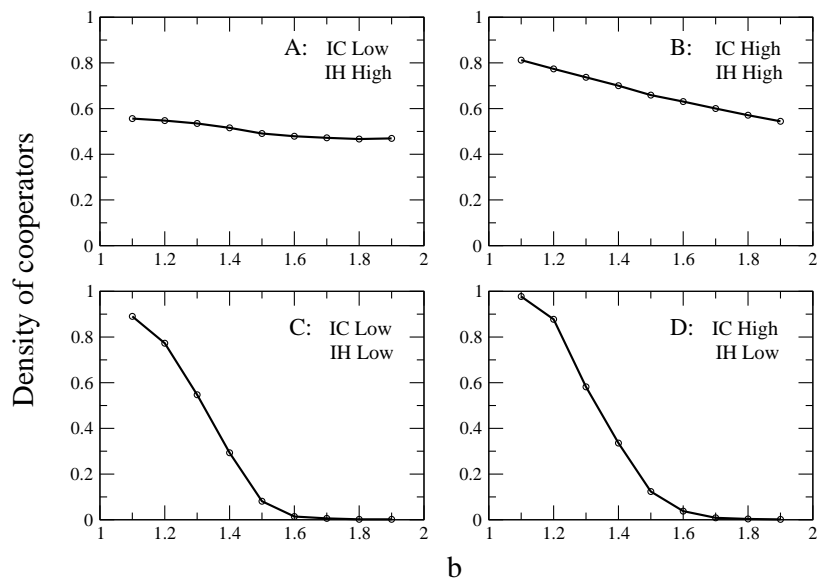


Fig. 3 Evolution of cooperation in four synthetic networks. Cases *A* and *D* correspond, respectively, to the synthetic versions of the PGP and email empirical networks. In case *A* communities were constructed as independent scale-free networks (Barabasi-Albert with $k_0 = 3$), and after they were sparsely interconnected with ($p_{inter} = 1.5 \times 10^{-5}$). Communities in case *D* were built as Erdos-Renyi random graphs ($p_{intra} = 1.5 \times 10^{-1}$), and the probability of interconnection between communities (p_{inter}) is 5×10^{-2} . Cases *B* and *C* were obtained from the former ones by increasing or reducing the probability p_{inter} . See [5] for more details on networks construction and simulation performance, and also a deeper analysis of results.

3 Resilience of cooperation

As we have seen above, from our previous work we have established that differences at a mesoscopic level among the two empirical networks influence their behaviour in terms of the evolutionary emergence of cooperation. In fact, the model we have developed to build networks with the desired values of IH and IC exhibits the same behavior and, therefore, can be used to design collaboration networks that are either very cooperative for low values of T (the incentive to forgo cooperation) or cooperative for large values of T or even try to improve performance in both respects. While this is generically interesting for the design of organization, one particular context where the application of these ideas can be very fruitful is the management of innovation networks, such as, e.g., the case of open innovation [13]. Therefore, it is worth to deepen on the role of the mesostructure on this issue. Recent studies have already pointed out the impact of the community structure on information transfer [14], there the counter-intuitive observation is that networks with

fuzzy community structure are more efficient in terms of information transfer than those with pronounced community structure. Our purpose in this paper is to report on new simulation experiments in order to assess the resilience of our two social networks and, subsequently, the resilience dependence on our two main control parameters, i.e., IH and IC.

In the literature about complex networks, the concept of resilience has been usually associated to the network robustness against the removal of their nodes. To be more precise, resilience has been understood as the capability to preserve large connectivity levels after those removals [3, 15, 16]. Consequently, the way to quantify the resilience in that line of research is to remove nodes (and their links) progressively, either at random (simulating failures of the connectivity system) or selecting them according to certain criterion such as their betweenness centrality, a measure of the number of shortest paths traversing a node, (simulating a purposeful attack on the network), and calculating the size of the largest connected component after the removal. In this work, following the wider understanding of the relationship among resilience and structure discussed in [17], we change the focus of our interest from the connectivity robustness to the quantification and comparison of the resilience of different organizations (represented by our two empirical social networks) against the (exogenous) introduction of defection.

The methodology used to perform the experiment, partially inspired on the one described above, consists of *attacking* both networks by introducing a certain quantity of defectors in a completely cooperative population, i.e., with a maximum cooperation level. Subsequently, the system is let to evolve according to the evolutionary dynamics discussed in Sec. 2.1, namely synchronous evolution with strategy update through unconditional imitation. Evolution proceeds until the system reaches a stationary state, where we measure the final density of cooperators.

A criterion for choosing the nodes to be attacked, i.e., those that will be occupied by a defector, must be provided. For our present experiment we have chosen their betweenness centrality. The reason is that this specific characteristic of nodes involves global knowledge about the structure of the network, and is of fundamental importance to detect bottlenecks in any diffusion process. Translating this experiment to a real life example, this *defection attack* could be seen as an attempt to destabilize an organization (a company, a public office, etc.) by bribing, blackmailing, or otherwise corrupting certain individuals occupying key positions. From this point of view, our choice of the betweenness centrality as our criterion for selecting nodes to be attacked is easily understood in terms of organizational issues such as the communication among groups [18] and the diffusion and control of information in organizations [19]. In the context of the application to innovation networks, these are of course important factors which would point wrongdoers towards the largest betweenness nodes.

3.1 Results

In Fig. 4, we show the results of applying the methodology proposed in the previous section to our empirical social networks. We can clearly observe that the PGP network is much more resilient to the introduction of defectors than the email one. As can be seen from Fig. 5, for the PGP network a successful attack needs, at least, to convert to defection a 25% of the whole population, the top quartile of the population ordered by decreasing betweenness. On the contrary, only one defector located in the most central position can lead the whole email network to defection. Recalling that the global characteristics of both networks are similar, we are led once again to conclude that their community structure is the responsible for such a dramatic difference in behavior.

A closer look to the results of simulations on the email network reveals an additional interesting phenomenon. For each value of N_D (number of introduced defectors), there is a minimum value of the incentive to defect, b , beyond which the final density of cooperators drops down to 0. This is also the case with the PGP network but, as already stated, for much larger values of N_D . The interesting point about the email network is that minimum b does not always decrease with increasing numbers of attacked nodes, but rather it oscillates until a value of N_D is reached such that a final full defection is reached for any value of b .

Building on our knowledge of the structural differences among the two networks at a mesoscopic level, we can advance an explanation of the observed behaviors and their differences. The community organization of the PGP network (weakly interconnected and internally stable due to the degree heterogeneity, see Fig. 1), restricts defection diffusion to a strictly local range (basically, the community where the defection is initiated). This explains why the final density of cooperators diminishes gradually. This behaviour only changes when the accumulation of small areas affected by defection affects a population large enough to saturate the whole system. On the contrary, the mesoscopic structure of the email network makes easier the diffusion of defection and, therefore, turns the whole network extremely sensitive to defection attacks. However, the picture is not so simple in so far as the most central positions of the email network are 'closer' (in terms of the number of nodes crossed by the shortest path among them) precisely due to the community structure. This circumstance has two consequences that can limit the spreading of defection through the network for a certain initial quantity of defectors. First, if two defectors are direct neighbours they decrease each others' payoff, reducing their probability to expand the infection among cooperator neighbours or, even, making them 'change their minds' and switch back to cooperative behavior. Second, as the number of initial defectors increases, it is more likely that two of them belong to the same community and diffuse defector behaviours across the same population of cooperator agents, reducing the efficiency of the process in terms of the ratio number of final defectors over number of initial defectors.




Fig. 4 Cooperation resilience in our empirical social networks when nodes are attacked as a function of their betweenness. Final density of cooperators as a function of b for different initial conditions, calculated over the email (top) and PGP (bottom) networks. The equilibrium densities of cooperators have been obtained by averaging 250 generations, after a transient time of 250 generation steps. Each point corresponds to an average over 500 independent simulations with the indicated number of defectors (in a completely cooperative population) as the initial condition.

3.2 Analysis using a toy model

In order to gain some insight on the above ideas about the influence of the community structure on the diffusion of defection and the phenomenon of ‘interference’ among defectors, we have performed a simple analysis by means of a toy model, that in spite of its simplicity captures the main observed features. The model is depicted in Fig. 6. Our calculation proceeds along the following lines: we consider different initial numbers of attacked nodes and for each one of them we calculate analytically the resulting effect (in terms of the final density of cooperators). As our update strategy rule in unconditional imitation, which is purely deterministic, the fate of our toy network can be easily predicted. The results are as follows:

A single defector at node 1. At the initial time step, the payoff for the defector is $3 \times b$ whereas the defector’s neighbors receive a payoff of 3. At the




Fig. 5 Cooperation resilience in the PGP network when a large percentage of nodes are attacked as a function of their betweenness. Final density of cooperators as a function of b for different initial conditions, obtained by averaging 250 generations, after a transient time of 250 generation steps. Each point corresponds to an average over 500 independent simulations with the indicated number of defectors (in a completely cooperative population) as the initial condition. PGP's high resilience is even more clear in this figure, since a successful attack needs to convert to defection, at least, the top quartile of the population ordered by decreasing betweenness.

subsequent step nodes 2, 3 and 4 switch to defection and spread defection to their communities for any $b > 1$.

Two defectors at nodes 1 and 2. At the initial time step, we have that $\text{payoff}_1 = 2 \times b$, $\text{payoff}_3 = \text{payoff}_4 = 3$ and $\text{payoff}_2 = 3 \times b$. With these values, at the next time step we will have a cooperator density of $9/13$, $8/13$ or $6/13$ depending on whether $b < 1.5$, $b = 1.5$ or $b > 1.5$, respectively. One further time step leads to cooperator densities of $10/13$, $8/13$ and 0 in the same cases. This implies a threshold of $b = 1.5$ beyond which full defection is the final state.

The same simple calculations lead to conclude that when 3 defectors are placed at nodes 1, 2 and 3 the final cooperator density is $5/13$ for any value of b , whereas if all four nodes are occupied by defectors full defection is always reached. In Fig. 7, we show all these results in a graphical way to facilitate the comparison of different cases. We observe that, despite the simplicity of




Fig. 6 Structure of the toy model used to study influence of the structure on the diffusion of defection

the structure of the toy model, it presents a counter-intuitive cooperation behaviour. While one defector at the one most central position is enough to reach a plenty defection, one additional defector can lead to high levels of cooperation (for low b values), and 2 more initial defectors are needed to get the full defection again. This clearly resembles our observation on the email network. The case of the PGP network can not be treated this way because, once again due to its community structure, a large fraction of the population must be changed to defection for the attack to succeed.

4 Discussion and conclusions

In this work we have considered the issue of the key role played by the community structure on cooperation in populations from the viewpoint of the resilience of the behavior. Our starting point was our recent work [5], where we demonstrate that two mesoscopic structural properties (the connectivity between communities, IC, and their internal structure, IH) condition the emergence and stability of cooperation in social networks. Here we report new simulation experiments and analytic approaches, aimed at studying in greater detail how IC and IH influence the cooperation robustness of social networks, and have interpreted the results in terms of organization's resilience to 'defection attacks'. We have adapted the concept of robustness or resilience to the context of cooperation in human societies by dealing with the ability of the social network under consideration to maintain its global cooperative behavior after key members stop cooperating and become defectors. Not surprisingly, we have found that these mesoscopic characteristics again turn out to be the




Fig. 7 Effect of introducing different amounts of defectors to the toy-model. Final density of cooperators as a function of b for the initial conditions studied analytically (see text). The counterintuitive behaviour presented by the email in Fig 4 is reproduced in this simple case. While attacks consisting on 1 or 4 defectors reach full defection, intermediate cases (2 and 3 defectors) lead to more cooperative scenarios.

crucial ones governing the response of the social networks to the appearance of defectors at relevant places; however, the specific results obtained for our two different cases are certainly interesting and deserve a discussion in depth.

In the case of the e-mail network, we are studying a social system defined by exchanges of e-mail among members of the different departments of an organization (specifically, a University). While some of those e-mail exchanges may relate to non-work contacts, we can safely assume that a large majority of them deals with work issues. However, the social network is not imposed, is self-generated and in fact, as shown in [6], the corresponding communities do not exactly match the externally organized departments. Such a network turns out to have large IC, i.e., communities are very well connected, and small IH, implying that they do not have clear leaders (hubs). Our conclusion in [5] that such a network behaves very well from the collaborative viewpoint but only for low values of the incentive to defect translates now, via our analysis of the resilience, into the conclusion that well directed attacks to a very small group of selected individuals are able to stop cooperation completely. It is important to realize that this dramatic effect takes place only for large enough values of the incentive to defect; however, when the attack affects to a slightly larger number of individuals, modifying only a 0.5% of the population leads to full defection for any value of b .

For our criterion to select individuals to attack, we have chosen betweenness centrality because, when studying topological resilience (i.e., the resistance of the network to become disconnected) those are the nodes that one should remove, as they are the ones intervening in more shortest paths. Here, where the network is not disconnected but it is the strategies of the neighbors that are affected, betweenness reveals itself again as a key property, in so far as it allows for a rapid propagation of the non-cooperative behavior (which is of course the goal of the attacker). It must be kept in mind, though, that knowledge of betweenness centrality requires in turn data on the whole network, as it is associated to a global feature (shortest paths between nodes). However, the attacker may overcome this problem by using node degree as a proxy for betweenness: Indeed, these two magnitudes are usually highly correlated (see [23] for a review of works comparing different centrality indicators), whereas knowledge of the degree requires only local information. We have repeated our simulations choosing degree as the attack criterion and we have found very similar results to those arising from betweenness. Therefore, we must conclude that such an attack is in principle possible and, as a consequence, that social networks with low IH and large IC are in general very vulnerable to disruption of the cooperative behavior by forcing a negligible number of individuals to defect. Interestingly, we have also shown that the attack may be less effective with a larger number of attacked nodes also due to the structural properties of the network, which shows that the evolution subsequent to the attack is by no means trivial.

On the contrary, the PGP network has proven to be very resilient to this kind of attack. PGP networks, being also of an undoubtedly social character, are different from the email case in so far as it is an informal organization from the very beginning, as it arises from spontaneous exchanges of passwords and encryption information between individuals, particularly at specialized conferences. As a result, the network is characterized by loosely connected communities with clearly differentiated hubs, which become so probably due to a preferential attachment mechanism acting when relations are established. As we have seen, this structure supports very resilient behavior, and a considerable fraction of the individuals (over a quarter of the population) must be corrupted to break down the cooperative behavior of the group as a whole. This conclusion is clearly related to the fact that scale-free networks have been claimed to be the best suited structure to support cooperative behavior [20]. Indeed, Poncela *et al.* [21] have recently studied the dependence of the observed cooperative behavior on scale free networks on the initial conditions and found that high levels of cooperation are the generic outcome of evolution, which is a clear indication of the robustness of cooperation on those systems. In this context, our results go beyond the work on model networks in a threefold way: We use real social networks, we have in mind the role of the mesoscopic structure (absent in the models) and we study resilience against attacking selecting individuals and not against random initial conditions.

Finally, our work must be placed in the wider context of organization design. While it is clear that ours is still a very simplified model, in so far as

the Prisoner's Dilemma is a particular social dilemma and interactions among individuals are much richer, our results point to a number of caveats that may prove useful for choosing a resilient setup for a group of individuals that must cooperate towards a common goal. Our main conclusion, namely that low IC and large IH are the best 'firewalls' against attacks, allows to use the model proposed in [5] to generate organizational charts tailored to show resilient cooperation. In this regard, it is interesting to realize that in our two examples, the one with the externally imposed social structure (email) performs much worse in terms of resistance to large incentives to defect and resilient cooperation than the self-organized network (PGP). We think it is not a coincidence that networks similar to PGP, with low IC and large IH, i.e., with a somewhat pyramidal structure, are found among large-scale criminal or terrorist organizations [22]. On a brighter note, our conclusion may also apply to open innovation networks [13], that arise, at least in part, from a volunteer participation mechanism that can yield a global structure similar to that of the PGP network. If that is the case, our study implies that the innovative collaboration will be resilient to the presence of individuals that free-ride on the true contributors to collective intelligence.

To conclude, we have shown that knowledge of the mesoscopic structure of a given social network is crucial to understand the resilience of the behavior of the individuals interacting on it. Resilience can be practically zero (such as on the email network) or very large (such as on the PGP network). As a consequence, mesoscopic features must be kept in mind when interpreting observations on social networks or designing specific-purpose networks. Our work opens the way to a more detailed study of this issue, which should proceed by considering other dilemmas or social interactions, as well as other examples of social networks.

References

1. Nowak MA (2006). Five rules for the evolution of cooperation. *Science* 314: 1560-1563
2. Nowak MA, RM May (1992). Evolutionary games and spatial chaos. *Nature* 359: 826-829
3. The structure and function of complex networks, M. E. J. Newman, *SIAM Review* 45, 167-256 (2003).
4. György Szabó and Gábor Fáy: Evolutionary games on graphs. *Phys. Rep.* 446 (2007) 97-216.
5. Lozano S, Arenas A, Sánchez A (2008) Mesoscopic Structure Conditions the Emergence of Cooperation on Social Networks. *PLoS ONE* 3(4): e1892. doi:10.1371/journal.pone.0001892.
6. Guimerá R, Danon L, Díaz-Guilera A, Giralt F and Arenas A (2003) Self-similar community structure in a network of human interactions. *Phys. Rev. E* 68: 065103.
7. Boguñá M, Pastor-Satorras R, Díaz-Guilera A and Arenas A (2004) Models of social networks based on social distance attachment. *Phys. Rev. E* 70: 056122.
8. Axelrod R, Hamilton W D(1981). The evolution of cooperation. *Science* **211**, 1390–1396.
9. Axelrod R (1984) *The Evolution of Cooperation*. Penguin, London.
10. Camerer C (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, NJ.
11. Crespi B J (2001) The evolution of social behavior in microorganisms. *Trends Ecol. Evol.* **16**, 178–183.

-
12. Duch J and Arenas A (2005) Community detection in complex networks using extremal optimization. *Phys. Rev. E* **72**: 027104.
 13. Chesbrough H (2003) *Open Innovation: The New Imperative for Creating and Profiting from Technology*. Harvard Business School Press, Boston.
 14. Danon L, Arenas A, Díaz-Guilera A (2008) Impact of community structure on information transfer. *Phys. Rev. E* **77**, 036103.
 15. Albert R, Jeong H, Barabási A-L (2000). Error and attack tolerance of complex networks. *Nature*, **406**, 378-382.
 16. Fontoura Costa L (2004) Reinforcing the Resilience of Complex Networks. *Phys. Rev. E*, **69**, (066127).
 17. Lozano, S and Arenas, A (2007). A Model to Test How Diversity Affects Resilience in Regional Innovation Networks. *J. Artif. Soc. Social Sim.* 10(4)8 <http://jasss.soc.surrey.ac.uk/10/4/8.html>.
 18. Granovetter, M (1973) The Strength of Weak Ties. *Am. J. Soc.* **78**, 13601380.
 19. Burt, R. S. (1995) *Structural Holes: The Social Structure of Competition*. Harvard University Press, paper back edition.
 20. Santos, F. C., Pacheco, J. M. & Lenaerts T. (2006) Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl. Acad. Sci.* **103**, 3490-3494.
 21. Poncela J, Gardenes J.G. , Floria L. M, and Moreno Y (2007) Robustness of Cooperation in the Evolutionary Prisoner's Dilemma on Complex Networks. *New J. Phys.* **9**, 184.
 22. Kaza, S., Xu, J., and Chen, H. (2005) Topological Analysis of Criminal Activity Networks in Multiple Jurisdictions,. In *The National Conference on Digital Government Research*, Atlanta, GA.
 23. Wasserman S and K Faust (1994) *Social Network Analysis: Methods and Applications*. Cambridge University, Cambridge.