new ground
Journal for Advances in Science

Research article | Published November 22, 2022 | Original research published in Proceedings of the National Academy of Sciences

Don’t Pack and Leave Too Soon: Forgiving Strategies of Cooperators Towards Defectors Can Lead to Cooperative Safe Havens

Cooperative behavior in biological networks is discouraged if defectors reap shared benefits. An obvious counterstrategy used by individuals is to avoid defectors by moving into another habitat. For spatially arranged ecological communities, mathematical analysis of the so-called snowdrift game shows under which circumstances this strategy has beneficial outcomes.

By Gil Henriques

Original research: Fahimipour, A. K., Zeng, F., Homer, M., Traulsen, A., Levin, S. A., & Gross, T. (2022). Sharp thresholds limit the benefit of defector avoidance in cooperation on networks. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(33), e2120120119. https://doi.org/10.1073/pnas.2120120119

Subjects: Cooperation, Metacommunity, Network, Game, Dispersal

New Ground article reviewed by: Thilo Gross

Cooperative behaviors – those that benefit others at a cost to the actor – are widespread in nature. Some mammals, including vampire bats and rats, often share some of their own food with hungry group members. Many microbes release valuable compounds into the extracellular environment, where they benefit nearby individuals. In the case of yeast, for instance, this is a sugar-digesting enzyme called invertase. However, a defector – say, a yeast cell that does not produce any invertase – still reaps the benefits provided by its more cooperative neighbors. It seems intuitive, then, that cooperators would have an incentive to abandon places in which the other individuals are not doing their share. Leaving in search of a better, more cooperative group would appear to benefit the migrants while simultaneously punishing the defectors who are left behind, depriving them of undeserved benefits.

But what would be the long-term consequences of this defector avoidance behavior? Would it favor more cooperative outcomes? In their recent PNAS publication “Sharp thresholds limit the benefit of defector avoidance in cooperation on networks,” Thilo Gross and his team present a mathematical model of the dynamics of cooperation in spatial settings and use it to study the consequences of defector avoidance strategies. One of their major findings – which may seem counterintuitive to some – is that immediate responses to defection do not appear to be the best choice.

Using games on graphs to analyze evolving populations as dynamical systems

A common approach used to study the evolution of cooperation is to model an evolving population as a dynamical system: a system of differential equations, each of which tracks the change over time in the frequency of occurrence for individuals of a specific type. In the simplest scenario, there are just two types: cooperators and defectors. When individuals interact, cooperators provide some sort of valuable collective benefit, whereas defectors do not. The benefit is shared equally with every individual in the interaction, even with defectors. Hence, it always pays to be a defector when interacting with a cooperator. However, when interacting with a defector, it might sometimes pay to be a cooperator, provided that whatever remains of the benefit after sharing outweighs the cost of the investment. When the model parameters are such that this is the case, it is referred to as a snowdrift game, and results in a state of equilibrium where a certain proportion of individuals cooperate and all others defect.

Regardless of the specific cost-benefit parameters, the question of what mechanisms can promote cooperation continues to drive research. One of the most powerful engines for cooperation is reciprocity: if individuals are more likely to interact with others of the same type, the benefits of cooperation preferentially flow toward cooperators. Reciprocity can occur, for example, when individuals occupy specific positions in space, so that they can only interact with their neighbors. In particular, if cooperators cluster together, they receive the benefits of mutual aid, with exploitation by defectors being restricted to the cluster's periphery.

This setup is often represented in the form of networks, or graphs, where each node is an individual and the edges represent spatial connections between them. These so-called "games on graphs" can be analyzed mathematically by making some additional assumptions (such as weak natural selection, i.e., that a single interaction has only a modest impact on an individual's success). Cooperators can then survive by grouping together into clusters. These clusters are "cooperative safe havens": regions of the network within which cooperators thrive.

Using metacommunity networks to investigate movement between habitat patches

This framework has been fruitful, but as Thilo Gross and his colleagues point out, representing physical space implicitly by means of a social network is highly abstracted from real-world ecology. This is because the network edges represent social rather than spatial connectivity. It is not clear, for example, how migration could be handled in such a setup. A different approach to the study of interactions in space – with deep roots in theoretical ecology – is metacommunity theory. Here, the population is divided into multiple habitat patches. Some patches are connected, so that individuals can migrate between them. This metacommunity can also be represented by a network, where each node represents a habitat patch containing many individuals. The subpopulation in each patch grows and changes over time, as its individuals interact with each other. Patches between which migration is possible are connected by edges. Such a metacommunity network explicitly represents, like a map, the physical space in which individuals live and move.

Studying interactions in metacommunities opens the possibility of investigating movement. Cooperators who find themselves surrounded by defectors could emigrate to some other location where they wouldn't be exploited, which would intuitively seem to benefit them. However, the movement of individuals tends to make different patches more homogeneous, meaning that, over time, all patches will display similar proportions of cooperators and defectors. This is the very opposite of a spatially segregated structure in which cooperators cluster into "safe havens," allowing reciprocity to thrive.

So, when should individuals stay and when should they leave in search of greener pastures? How do different defector avoidance strategies – strategies such as "leave as soon as someone defects against you" or "leave only after interacting with defectors three times in a row" – affect the maintenance of cooperative safe havens?

A sharp threshold in connectivity governs the development of either heterogeneous or homogeneous states

As an illustrative example, Thilo Gross and colleagues present a simple network with only two patches. Each patch is home to a subpopulation consisting of cooperators and defectors engaged in a snowdrift game. A system of two differential equations describes the change over time in the frequency of each type. In the simplest possible case – when there is no migration between patches – these changes result from the balance between the reproduction and mortality rates within a given patch. The reproduction rate of an individual is proportional to the payoff that it achieves during its interactions, whereas the mortality rate is proportional to the density of the respective patch's population (a common modeling choice that accounts for competition for limited resources and prevents the population from growing infinitely). In the absence of migration, both patches evolve toward the same equilibrium state, where cooperators and defectors coexist at a certain frequency (which depends on the game's parameters) and in a homogeneous state, where the two patches are home to identical subpopulations.

What if we allowed for migration between the patches? Then, changes in the number of cooperators and defectors would no longer be due to reproduction alone, but also to migration between patches. The actual rate of migration would depend on the strength of connectivity between the patches as well as on the dispersal strategy: the rules that determine when an individual decides to leave a patch. For this example, Gross and his team present a straightforward dispersal strategy where cooperators leave a patch if they have been cheated several consecutive times. Mathematically, this dispersal strategy implies that the rate of emigration increases more than linearly with the proportion of defectors. The more forgiving cooperators are – that is, the more defections they tolerate before deciding to leave – , the stronger this nonlinearity is.

If the patches are weakly connected, such that overall migration is rare, the outcome is the same as before: the system evolves to the same homogeneous stable state, with each patch characterized by the same relative frequencies of cooperators and defectors. Increasing the connectivity does not affect this equilibrium until a critical value is exceeded. Above this value, the system stabilizes at a completely different, heterogeneous equilibrium, where cooperators dominate one patch and defectors the other. Below this value, the outcome will be homogeneous. Emigration to avoid defectors does not benefit cooperators unless this sharp threshold is crossed.

A master stability function decouples effects of the game from spatial structure

After having presented this simple case, the authors introduce the general theory, which allows for a wide variety of scenarios, including complex network shapes, games with more than two types, more complicated decision rules for when to stay or leave the patch, and strong natural selection (there are a few assumptions, including that the patches must not differ in the quality of habitat). The basic mathematics are similar: a given type of individual within a patch still increases in frequency via reproduction and immigration and decreases via mortality and emigration. Individuals can migrate between patches that are connected, with rates of migration that depend on the connection strength and on their dispersal strategies. The main challenge is that there are many more variables to keep track of over time: one for each type of individual in each patch.

In any dynamical system, we are usually interested in investigating the stability of different outcomes. For instance, given the two-patches example in the previous section, we are interested in asking whether the homogeneous outcome is stable (which was the case at low migration frequencies) or whether, instead, a heterogeneous outcome will be reached, with some patches dominated by cooperators and others dominated by defectors (which was the case at high migration frequencies). To determine stability, it is necessary to compute the eigenvalues of a matrix of derivatives called a Jacobian, which contains information about how the increase or decrease in the frequency of each type of individual (in each patch) is accelerated or decelerated by changes in the composition of the population.

Given the general model's complexity, this might seem like a Herculean task. However, it turns out that even here, stability can be studied analytically. This is because, in the metacommunity network framework, it is possible to calculate what is known as a master stability function. This approach breaks down the large, complicated Jacobian matrix of the full network-structured system into simpler terms. The first term describes the effects of the interactions (the game which is played within each patch): it is based on the much simpler and smaller Jacobian of the in-patch subpopulation. The second term encodes the effects of the network's shape: it is based on two matrices, which together describe the connections between the patches and how the population sizes within a patch accelerate or decelerate the change over time in the migration rates.

In summary, the master stability function disentangles the influence of the specific spatial structure from that of local interactions, i.e., from the rules and parameters of the game. The spatial structure determines the eigenvalue that characterizes the fate of the system: if this eigenvalue exceeds a critical threshold, the system's equilibrium becomes heterogeneous. The critical threshold itself depends on the parameters governing local interactions.

Forgiveness generates safe havens for cooperation

What then is the effect of movement? As previously mentioned, the movement of individuals tends to make different patches more homogeneous. However, the two-patches example and the master stability function approach both show that as connectivity strength surpasses a certain threshold, the community can become spatially segregated, with the formation of cooperative safe havens.

As the authors show using numerical simulations, strong connectivity only leads to the formation of safe havens when the dispersal strategy is such that emigration increases more than linearly with the percentage of defectors. A strategy where cooperators emigrate as soon as they have been cheated (which results in an emigration rate that is linearly proportional to the proportion of defectors) would simply lead to too much mixing between patches, contributing to a completely homogeneous network. In contrast, when cooperators are so forgiving that they only leave after experiencing several consecutive defections, then safe havens for cooperation may result. The more forgiving cooperators are, the less network connectivity strength is required to generate spatial heterogeneity.

But what do these spatially heterogeneous systems look like? The authors simulated randomly generated networks, for which the connectivity strength exceeded the threshold value. Just like in the example, cooperators and defectors became spatially segregated, with the formation of safe havens within which cooperators interacted frequently with one another, reaping the benefits of mutual aid. Furthermore, the cooperators congregated in well-connected patches (that is, easily accessible by migration).

These simulation results are based on a specific type of dispersal strategy: those in which the cooperators leave a patch whenever they have been consecutively cheated a certain number of times. Of course, many other strategies are conceivable. To study more general dispersal strategies, the authors exploited a similarity between the models where spatial structure is represented using metacommunity networks, and a different family of models, where it is represented using a continuous physical space (for instance, a plane defined by two coordinates, such as latitude and longitude).

In continuous space, heterogeneous patterns can emerge spontaneously from initially homogeneous, uniform conditions. By relying on a strong analogy between the master stability function and the equation governing pattern formation in continuous space – the Turing instability in partial differential equations –, the authors found that, in the highly connected systems where heterogeneities are most likely, the emergence of safe havens is impacted more by the choice of dispersal strategy than by the local game dynamics. In particular, the formation of spatially heterogeneous networks is fostered by dispersal strategies where cooperators vigorously suppress the emigration of other cooperators in their own patch. The forgiving strategy discussed above is an example of this condition since, under said strategy, cooperators are less likely to emigrate when they are surrounded by many fellow cooperators.

Future directions

The generality of the metacommunity model introduced in the article opens many prospects for future work. For instance, although the authors focused on snowdrift games with two types of individuals, the model can be applied to various kinds of interaction. The literature is full of games with large catalogs of interacting types, with fanciful names like "reciprocator," "loner" or "punisher." In some of these games, the frequencies of the interaction types can cycle over time, rather than stabilizing at some intermediate value. How would these scenarios play out in metacommunity models?

Another avenue that can be explored using the model is the effect of adaptive dispersal strategies: what would happen if individuals chose to migrate to habitats that were already populated by cooperators, rather than moving blindly to a patch they knew nothing about? Would this allow defectors to break into the cooperative safe havens?

Alternatively, future work could expand the model by relaxing some of its assumptions to account for more realistic scenarios. For example, resources in nature are heterogeneously distributed across space. What would happen if the network patches had habitats of different quality (or if their quality changed over time)?

Conclusion

Gross's team introduces a method for performing stability analysis in metacommunity models of cooperation games. Using the master stability function, they can explore whether the system evolves toward spatial homogeneity (where each patch is home to the same proportion of defectors) or whether cooperators segregate into safe havens (where they can benefit from mutual rewards).

Perhaps counterintuitively, they find that selective emigration of cooperators away from defectors does not always lead to greater benefits for cooperators. Only when connectivity between patches is strong enough to surpass a threshold can safe havens form. Furthermore, this outcome depends on dispersal strategies where cooperator emigration increases more than linearly with the percentage of defectors in the patch.

Among these are forgiving strategies, in which cooperators do not abandon the patch as soon as they're cheated. By showing some leniency before walking away, cooperators contribute to the formation of spatial heterogeneity. Having done so, when they finally decide to pack up and leave, they are more likely to actually find a safe haven. •

The researchers

Thilo Gross is a network and data scientist at the Helmholtz Institute for Functional Marine Biodiversity and a professor of biodiversity theory at the University of Oldenburg, Germany. Ashkaan K. Fahimipour is an assistant professor at the Department of Biological Sciences, Florida Atlantic University.

Read more

Defector Avoidance in cooperation games on networks
On YouTube, Thilo Gross presents a 12-minute video abstract of his paper.