Multilayer Networks and the Evolution of Risky Cooperation

Matthew Coates*

Abstract

Philosophers have shown that social networks significantly influence the emergence of prosocial behaviors. However, despite real social communities often being characterized by agents acting in different social spaces at the same time, network structures that capture this, such as multilayer networks, have received little study. In this paper, using the stag hunt game, I show that multilayer networks have a significant effect on the extent to which cooperation emerges in cases where cooperation is risky. Given many real social communities are multilayer, I therefore argue that when studying the impact of networks on prosocial behaviors, multilayer networks should be investigated.

1 Introduction

Philosophers have shown that social ties and social networks significantly influence the cultural evolution and emergence of prosocial behaviors (Skyrms, 2004; Zollman, 2005;

*Department of Logic and Philosophy of Science, UC Irvine.

Alexander, 2007). Altruism, cooperation and bargaining all evolve differently on social networks. Given the importance of social networks to human life, these results have deep implications for human moral behavior.

However, a group of network structures that have received little attention in formal philosophy are multilayer networks. Multilayer networks are a structure consisting of multiple layers where each layer contains a network, with interlayer networks consisting of connections between the nodes on different layers. This is despite real social communities often being characterized by agents acting in different social spaces at the same time. For example, individuals have differing social ties with family, friends, work collaborators and so on. It is realistic that individuals may distinguish their behavior based on the specific social space that they are currently interacting in. Multilayer networks provide a way to model such scenarios.

Whilst receiving little attention in philosophy, multilayer networks have been widely studied in other disciplines (Bianconi, 2018). Various studies have shown that interdependent multilayer structures can dramatically influence evolutionary dynamics in games like the Prisoner's Dilemma (Gomez-Gardenes et al, 2012).

The primary aim of this paper is to show that multilayer networks have a significant effect on the extent to which cooperation and trust evolve in cases where cooperation is risky, using the stag hunt game. In many single-layer network structures the spread of cooperation in the stag hunt is very rare, or even impossible. I demonstrate that compared to single-layer structures, multilayer structures can increase the likelihood of cooperation spreading in the stag hunt game under certain conditions, and decrease it under others. Additionally, I show that with multilayer networks we see the emergence of agents distinguishing cooperative behavior based on social space without assumptions about the meaning of those social spaces.

Using these results, I then claim that multilayer networks are relevant to philosophical investigations of the emergence of human moral behavior more generally. My results show multilayer network structures can significantly impact the dynamics of a network model. Given many real social communities are multilayer, I therefore argue that when studying the impact of social networks on cultural evolution, multilayer networks should be considered.

I begin by giving a basic introduction to multilayer networks, before describing my model. I then demonstrate the effects of multilayer network structures on the spread of cooperation in cyclical networks, a case it struggles with in single-layers. Next I consider the case of Erdos-Renyi random networks. I demonstrate through simulations that increasing the number of layers increases the likelihood of trust and cooperation spreading, and leads to a higher average trust and cooperation in the network.

2 Multilayer Networks

A multilayer network is a network structure formed by multiple interacting networks. It consists of M distinct layers, each containing a network called an intralayer network, as well as $M(M-1)/2$ networks which describe the interactions between nodes on each pair of layers, called interlayer networks. Multilayer networks capture that real systems are often formed by many, interconnected, networks, where interactions on each may take different forms.

There are two main types of multilayer networks (Finn et al, 2019):

- *Multiplex Networks* are multilayer networks where the same set of nodes appears in each layer. The interlayer edges connect nodes to themselves on different layers, and no node has an interlayer edge between itself and a different node. Each layer may then represent some different type of interaction, or the same type of interaction but at different time points where the network structure may be different (a *temporal network*).

- *Interconnected network* are multilayer networks where different sets of nodes appear in each layer. The nodes in each layer do not necessarily correspond to the same entity. The interlayer edges can then connect nodes representing different entities between layers. For example, a network with one layer of pollinators and one layer of plants.

Many animal communities are multilevel, and can be represented by multilayer networks (Finn et al, 2019). For example, during breeding season superb fairy-wrens live in communities consisting of breeding groups. However outside of breeding season, they form multilevel communities where one level consists of breeding groups, another consists of multiple breeding groups to form supergroups, and a final level consists of interactions with birds in other breeding groups and super groups. The birds interact in different ways with other birds depending on the layer they are connected in. These interactions may relate to raising infants, foraging for food and protection from predators (Camerlenghi et al, 2022).

Human societies are also often multilayer. For example, social networks between people may consist of many different social ties with different interactions and meanings, such as family, friends and work colleagues. As real human societies demonstrate this sort of structure, multilayer networks are important for understanding social interactions and the emergence of human prosocial behavior.

3 The Model

I now describe my model. The structure of this model is a multiplex network. There exists a set of layers, $L = \{1, \ldots, M\}$, with each layer consisting of a network made up on n nodes, where nodes represent agents. Each numbered node represents the same agent in each layer. For simplicity, I assume that the network structures are generated the same way in each layer. For example, for cyclical networks, every layer will be made up of cyclical networks. The connections between nodes in each layer are determined independently by the network generation; every agent is represented by one node in each of the M layers, however their neighbors may differ in each layer.

The stag hunt game is a widely used game for studying the emergence cooperation under risky conditions (Skyrms, 2001; Skyrms, 2004; Alexander, 2007; Huttegger and Smead, 2011). It has the payoff structure seen in Table 1. Each agent gets a guaranteed payoff if they hunt hare, but can get a higher payoff if they cooperate and hunt stag. However, they risk getting nothing if they hunt stag but their partner does not. This then relates to the emergence of trust, as two agents who trust each other to cooperate and hunt stag gain higher payoff.

At the start of each run, each agent is assigned a strategy (Stag or Hare) at random in each layer to use against their neighbors. They may be assigned different strategies in each layer.

For simplicity, I assume $y = z$. I also assume $y + z > x$, so that Hare is the risk dominant strategy. This is when Stag is riskier than hare, and therefore is most worth

Table 1: Stag Hunt Payoff Structure

	Stag	Hare
Stag	x,x	0,y
Hare	y,0	z,z

focusing on when studying the evolution of risky cooperation.

At each time step, each agent plays a stag hunt game against all of their neighbors in every layer. Agents will receive a *layer payoff* for each layer, which is the sum of their payoffs in the games they played on that layer. Each agent then has a *net payoff* which is the sum of all of their layer payoffs.

Just as in a single-layer network, the updating dynamics have a large effect on the results of the model. For example, if each agent could simply see each others layer payoffs, and update their strategy in that layer in light of that, then the dynamics would simply be the same as many single-layer networks.

Therefore, I model the updating as follows. After playing the games on each layer, the agent updates their strategy in one layer, chosen at random among the M possible layers. I assume each agent uses imitate-the-best learning (Alexander, 2007). The agent adopts the strategy from their immediate neighbor in the chosen layer who has the highest net payoff, if it is higher than their own. In the next round, they will use that new strategy in games that layer. This system of updating makes the dynamics on one layer dependent on what is happening on every other layer.

This represents a scenario where agents observe overall success of their neighbors, but not each interaction contributing to that success. An agent may know that their neighbor on the layer is more successful than them, but they do not know how their neighbors strategy on that layer helped contributed to their total success. However, because the agent wants to be more successful, they are willing to take the risk and copy their neighbor on that layer in case it did help that total success.

I illustrate with an example. Imagine two fairy-wrens, A and B, from different supergroups meet and interact. Afterwards, A notices that B has foraged more food. A decides it wants to be as successful as B, however it does not know whether B's total food was a result of interactions with other birds outside of its supergroup, its interactions with birds within its supergroup, or interactions with birds within its breeding group. It just knows that B is more successful at foraging food. A therefore decides to copy B's strategy in its interactions with birds outside of its supergroup, as that is the only information it has to do as well as B did.

4 Results

4.1 Cyclical Networks

I first consider cyclical network structures. As shown by Alexander (2007, 114), a similar model with one layer and Imitate-the-Best learning converges to a fixed state consisting of a mixture of Hare and Stag after one generation. To demonstrate this, consider the payoffs $x = 3$, $y = z = 2$. A Hare will always receive a payoff of 4, no matter who their neighbors are. A Stag surrounded by two Hare neighbors receives a payoff of 0, so will switch to Hare. Likewise, a Stag in a cluster of only two Stag (so both they, and their neighboring Stag, also have Hare neighbors) will receive a payoff of 3, so will still switch to Hare. Finally, clusters of 3 or more Stag are stable. The Stag on the boundary of the cluster still only receives a payoff of 3, but Stag's within the cluster receive a payoff of 6,

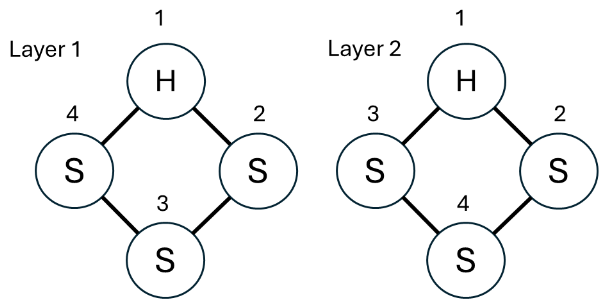

Figure 1: Configuration allowing Stag to spread in a 2 layer network.

so those on the boundary will not switch to Hare because they have a neighboring Stag with a higher payoff.

Therefore, we end in a fixed state where clusters of 3 or more Stag are stable, and everyone else is a Hare. The corollary of this is that it is impossible for Stag to spread in a single-layer cyclical network. A Hare will always receive a payoff of 4, but a Stag bordering a Hare will at most receive a payoff of 3 because they must have at least one Hare neighbor. This means no Hare will switch strategy.

Is this the same in the multilayer case? I show analytically it is possible for Stag to spread in a multiplex, cyclical, network, so that there are a greater number of Stag at the end state than there were in the beginning state:

Take a multiplex network of 2 layers. Each layer consists of 4 nodes, with the starting assignment of Stag and Hares as shown in Figure 1. After playing against their neighbors on each layer, Player 1 has net payoff of 8, Player 2 has net payoff of 6, and Players 3 and 4 have net payoff of 9. Let's say that for Player 1, Layer 1 is chosen to update on. They see that their neighbor, Player 4, has a higher net payoff than them, so they adopt Player 4's Stag strategy on Layer 1. In contrast, no Stag will switch because either they have a higher net payoff than the Hare, or they're connected to a Stag with a

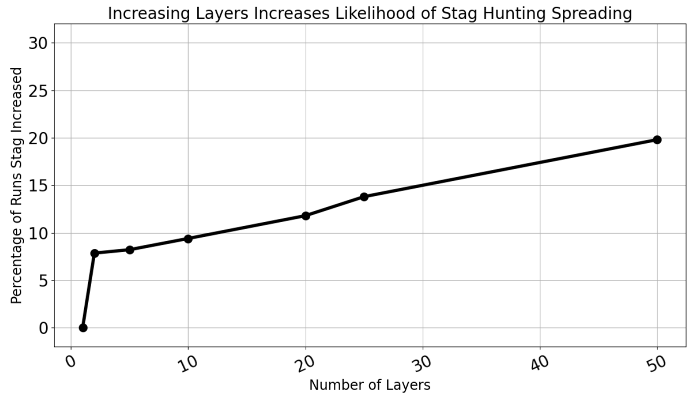

Figure 2: Percentage of runs where Stag increases with n = 5.

higher net payoff than the neighboring Hare. The same dynamics repeat when Player 1 updates in Layer 2. This means that the amount of players using Stag has increased in between rounds, showing that it is possible for Stag to spread in multiplex cyclical networks.

However, whilst it may be mathematically possible, this says nothing about whether it is likely. I use computational simulations to determine whether this is a likely scenario. Each simulation is run 1000 times. I begin by considering a small network with only 5 nodes on each layer. As seen in Figure 2, showing the percentage of runs where total usage of Stag increases between the start and end of the simulation, it is a relatively rare scenario for Stag to spread, even without considering how far it spread. Note that in this scenario, Stag could still spread to some agents, but since Hare is spreading elsewhere there is no global increase in Stag.

The spread of Stag is relatively rare because it requires a particular network setup. Take the 2 layer case. A Hare will receive a payoff of 8 across both layers (4 from each layer), no matter who their neighbors are on each layer. Therefore to switch to Stag they must have a neighbor who receives a payoff greater than 8. The only way for a Stag to receive a payoff greater than 8, whilst also being neighbors with a Hare on a layer, is to be neighbors with 2 other Stag on one layer (payoff of 6), and neighbors with a Stag and a Hare on the other layer (payoff of 3) to give a total payoff of 9. The Hare then needs to update on the layer on which they are neighbors with the Stag with payoff 9. To then spread again, it would need to put a further Hare into this same position and so on. It requires a very specific configuration of Stag and Hares to even spread once, never mind to continue spreading, whilst Hare will spread more easily for the same reason as in the single-layer case.

Whilst it remains a relatively rare scenario, as seen in Figure 2, as the number of layers increases there is an increase in the likelihood of Stag spreading. This is because increasing the number of layers increases the number of configurations that allow Stag to spread. This makes initial conditions allowing for Stag spreading more likely.

Less obviously, an increase in Stag also depends on network size, with larger networks on each layer being much less likely to see an increase in Stag. When run with 100 agents on each layer there are no cases where Stag increases. I tested this with up to 50 layers. This result is because Hare still spreads more easily than Stag. In smaller networks, when the configuration needed for Stag to spread occurs, it is less likely to have enough agents for Hare to also spread elsewhere on the network. In contrast, in larger networks these configurations still occur, there are enough agents that Hare will also likely spread somewhere else on the network and cancel out the increase in Stag.

These results suggest that cooperation is more likely to evolve when agents are taking part in small communities, but have interactions in many different social spaces and tailor their action in each. However, its likelihood of evolving is still relatively rare.

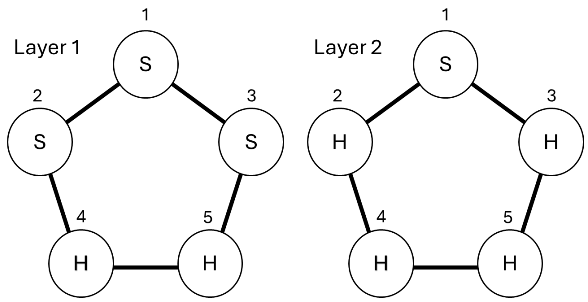

Figure 3: Configuration preventing Stag from spreading in a 2 layer network with update neighborhood larger than interaction neighborhood.

What if the update neighborhood is larger than the interaction neighborhood, so that agents still only play against their neighbors, but will update on the payoffs of agents connected in a longer path? In the single-layer case, when there is a cluster of 3 or more Stag in the initial state, Stag will spread and this will lead to an all Stag state. In contrast, with a larger update neighborhood multilayer networks actually do worse than single-layer networks, as the presence of a cluster of 3 Stag on one layer is not sufficient for Stag to spread.

To demonstrate, take a multiplex network of 2 layers with update neighborhood of 2. Each layer consists of 5 nodes, with starting assignment shown in Figure 3. After one round, Player 1 has a net payoff of 6, Players 2 and 3 have net payoff of 7, and Players 4 and 5 have net payoff of 9. Layer 1 is chosen to update on for Player 4 or 5. Though they can see Player 1 in the middle of a cluster of 3 Stags, they will not switch to Stag, because they already have a higher net payoff. Additionally, no matter which layer Player 1 updates on they will switch to Hare, as they can see Players 4 and 5 with greater payoff.

Nonetheless, based on simulations, Stag does usually spread on at least one layer of the network, with the end state likely consisting of a combination of Stag and Hare on each layer. The number of Stag and Hare on each layer typically differ, and one common end state is where Stag has spread through most of one layer, whilst Hare has spread on the other.

Though multiplex structures decrease the likelihood of full cooperation, this result highlights an important dynamic for the evolution of risky cooperation. We frequently see agents in the model evolve to trust and cooperate in some layers, whilst not cooperating in others. With 2 layers and $n = 50$, Stag increased on exactly one layer in 53.6% of runs. As layers increase, the likelihood of Stag both decreasing and increasing on at least one layer in the same run increases, reaching 79.1% with 3 layers. Real humans often adhere to community specific norms and conventions, choosing different actions based on the social space they are acting in. In this model we see the emergence of community specific norms for cooperation with minimal assumptions about the meaning of each social space. Therefore, this may capture something real about the emergence of prosocial behaviors.

4.2 Random Networks

Cyclical networks may not be the most representative of actual communities, therefore I also consider random networks. I use the same generation algorithm as Gomez-Gardenes et al (2012), where each layer consists of an Erdos-Renyi random graph, each with the same number of nodes and the same average degree $< k >$. I take the size of each layer to be n = 100, and I consider two cases. The first is a sparse graph with average degree $< k >= 3$, and a more connected graph, with average degree $< k >= 20$. I do this to

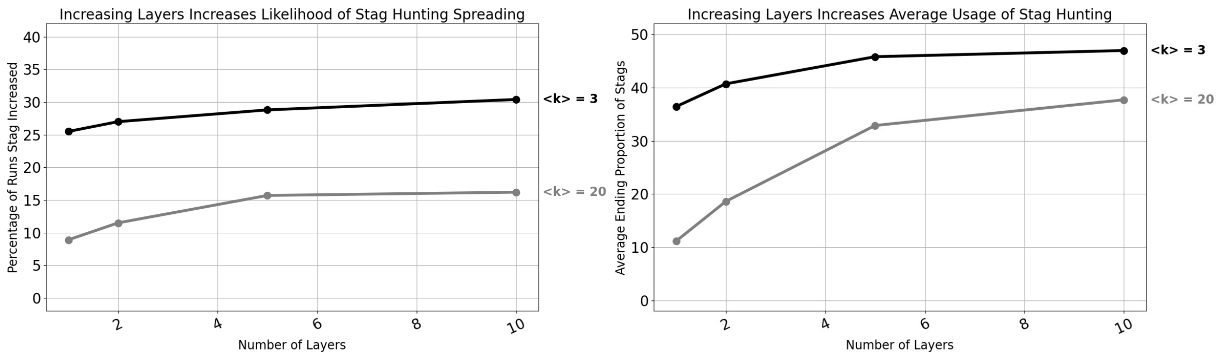

Figure 4: On a random network with n = 100 (a) Percentage of runs where Stag increases. (b) Average ending proportion of Stag usage.

test whether average degree effects the results. Finally, I run each simulation 1000 times, with 1000 rounds of updating in each run.

With random networks it is possible for Stag to spread in a single-layer network. The question then is whether multiplex networks improve the likelihood of Stag spreading. To analyze this I consider two different measures. The first is how often Stag spreads, in the same way as in the cyclical network case, the second is the average proportion of agents using Stag at the end of each run.

First, I consider the sparse network case, with average degree $< k >= 3$. The general result is that multiplex networks help Stag to spread. I find that as the number of layers increases, the more likely it is for Stag to spread, and the greater the average proportion of Stag at the end of the run. This can be seen in Figures 4(a) and (b).

However, despite these results I do find that a single-layer network can end with far more Stag than a multiplex network. This can be seen in Table 2. There are two reasons for this. The first is that a single-layer network consists of less nodes, so the probability of close to all being assigned Stag at the beginning is higher than when multiplex. The second is that with more layers the net payoffs for each individual become more even, meaning that results are less variable. Additionally, we see the same effect as in cyclical networks where a common scenario is the emergence of layer specific strategies, where Stag spreads on some layers and Hare spreads others. With 2 layers and $n = 100$, Stag increased on exactly one layer in 52.1% of runs.

Table 2: Final results for 1000 randomly generated networks with $< k >= 3$.

Proportion of Stag	1 Layer	5 Layers	10 Layers
p =1	0	0	0
$0.9 \leq p < 1$	46	0	0
$0.8 \leq p < 0.9$	88	0	0
$0.7 \leq p < 0.8$	51	3	0
$0.6 \leq p < 0.7$	39	54	11
$0.5 \leq p < 0.6$	63	268	300
$0.4 \leq p < 0.5$	83	434	574
$0.3 \leq p < 0.4$	115	199	113
$0.2 \leq p < 0.3$	158	40	2
$0.1 \leq p < 0.2$	171	2	0
$0 < p < 0.1$	186	0	0
p = 0	0	0	0

What effect does connectivity have? Looking at average degree $< k >= 20$, I find that degree has a large effect on the spread of Stag. Stag is both far less likely to spread when the average degree is higher and far more likely to be destroyed completely. However, I still find that multiplex networks are more likely to protect the existence of Stag in the network and are still more likely to allow it to spread. The larger the number of layers, the greater likelihood of Stag spreading, and the greater the average proportion of ending Stags. These results can also be seen in Figure 4.

Therefore, Stag is able to survive and spread more often in random multiplex networks than in random single-layer networks. Additionally, this increase in success for

Stag due to multiple layers even allows Stag to spread in cases where is would be difficult to spread in the single-layer case, such as in networks with high average degree. Finally, random multiplex networks still allow for the similar emergence of layer specific strategies to cyclical multiplex networks.

5 Conclusion

This paper has been motivated by the concern that the multilayer structure of real social communities has received little attention by philosophers when modeling the emergence of prosocial behaviors. This led to two main aims. The first was to study the effect that the multilayer character of social interactions has on the evolution of cooperation and trust under risky conditions, using the stag hunt game. The second, more general, aim was to show that multilayer network structures are relevant to philosophical investigations into the emergence of human moral behavior.

Regarding the first aim, previous results on single-layer networks have shown that network structure can play a significant role in whether cooperation emerges in the stag hunt (Skyrms, 2004; Zollman, 2005; Alexander, 2007). Under many conditions, the network structure works against cooperation, making its evolution rare or even impossible (Alexander, 2007, 114). I find that incorporating multilayer network structures provides ways for cooperation to spread in these circumstances. Additionally, even in some network structures where cooperation can already spread in single-layer networks, such as Erdos-Renyi random networks, cooperation is still promoted further in multilayer networks.

These results may suggest that a multilayer character of social interaction helps with the evolution of cooperation. However, under other conditions a multilayer structure hinders the spread of cooperation when it would otherwise spread on a single-layer network. Nonetheless, whilst cooperation is hindered under these conditions, the end state may better account for real practices we observe, and provide an explanation for the emergence of cooperation in real communities. In these cases, though cooperation is unlikely to spread to the entire population on both layers, it does often spread on at least one layer. A common end state is most agents cooperating in some layers but not in others. This reflects the fact that humans often adhere to community specific norms and conventions, changing behavior based on the social space they are acting in. These findings demonstrate that such community specific norms relating to cooperation can emerge with minimal assumptions about the meaning of specific social spaces.

These results lead into my second aim. I have shown that multilayer network structures do lead to differences in dynamics that may otherwise be missed by only using single-layer networks. Given real communities are often characterized by agents acting in different social spaces at the same time, distinguishing their behavior in each, these different dynamics may better capture the real phenomena that is being modeled. Therefore, when studying the impact of social networks on cultural evolution, multilayer networks should be utilized as they may increase the explanatory value of such models.

As a final point, we can clearly see that network structure has varied and surprising impacts on the dynamics of cultural evolution in models. Given how varied the impacts can be, we should be cautious about drawing quick and general conclusions from simple network models. Results purportedly about the evolution of human moral behavior which do not capture the complexity of real social networks may be misleading. Whilst multilayer networks may capture more of the complexity than other simple networks, even they likely do not capture all of it.

References

Alexander, J.McKenzie. 2007. *The Structural Evolution of Morality.* Cambridge: Cambridge University Press.

Bianconi, Ginestra. 2018. *Multilayer Networks: Structure and Function.* Oxford: Oxford University Press

Camerlenghi, Ettore., Alexandra McQueen, Kaspar Delhey, Carly N. Cook, Sjouke A. Kingma, Damien R. Farine, and Anne Peters. 2022. "Cooperative breeding and the emergence of multilevel societies in birds". *Ecology Letters*, 25(4), 766-777.

Finn, Kelly. R., Matthew J. Silk, Mason A. Porter, and Noa Pinter-Wollman. 2019. "The use of multilayer network analysis in animal behaviour". *Animal Behaviour*, 149, 7-22

Gomez-Gardenes, Jesús., Irene Reinares, Alex Arenas, and Luis Mario Floría. 2012. "Evolution of Cooperation in Multiplex Networks". *Scientiic Reports*, 2:620

Huttegger, Simon., and Rory Smead. 2011. "Efficient social contracts and group selection". *Biology and Philosophy* 26, 517–531

Skyrms, Brian. 2001. "The Stag Hunt". *Proceedings and Addresses of the American Philosophical Association*, 75(2), 31–41.

Skyrms, Brian. 2004. *The Stag Hunt and the Evolution of Social Structure.* Cambridge: Cambridge University Press

Zollman, Kevin J. S. 2005. "Talking to Neighbors: The Evolution of Regional Meaning." *Philosophy of Science* 72(1): 69–85.