

How the integration of group and individual level selection affects the evolution of cooperation

James Jahns⁴ and Arend Hintze^{1,2,3,4,*}

¹ Department of Integrative Biology

² Department of Computer Science and Engineering

³ BEACON Center for the Study of Evolution in Action

⁴ Michigan State University, 288 Farm Lane Rd, East Lansing, Michigan, USA

* hintze@msu.edu

Abstract

Many evolutionary models that explore the emergence of cooperation rely on either individual level selection or group level selection. However, natural systems are often more complex and selection is never just on the level of the individual or group alone. Here we explore how systems of collaborating agents evolve when selection is based on a mixture of group and individual performances. It has been suggested that under such situations free riders thrive and hamper evolution significantly. Here we show that free rider effects can almost be ignored. Sharing resources even with free riders benefits the evolution of cooperators, which in the long run is more beneficial than the short term cost.

Introduction

The tragedy of the commons (Hardin, 2009) describes a situation in which group members have to make a decision between investing into a common pool or withholding their contribution due to self-interest. Since evolution prefers individuals based on their advantage over others, selfishness is the expected outcome of selection in this game. The question then becomes, what factors would allow cooperation to evolve? Previous work found that cooperation can become the natural outcome of evolution under specific conditions (Adami et al., 2016). However, external factors such as kin selection, direct reciprocity, indirect reciprocity, network reciprocity, or group selection (Nowak, 2006) have been identified as external factors that promote cooperation. Simply put, any of these mechanisms provides enough benefit to overcome selfish behavior so that cooperative behavior can evolve.

One of these mechanisms is the theoretical concept of group selection (Wilson and Sober, 1994), where the group as a whole is under selection instead of the individual, and its genetic persistence is based on a group-level trait. This mechanism works perfectly well in social insects such as ants and bees, since here the reproductive success of the queen is directly dependent on the performance of the colony. However, in these examples, we also only find one member of the group (the queen) to have reproductive capabilities, which eliminates competition about reproductive success between

other members of the group (also known as an extended phenotype (Dawkins, 2016)). Most other organisms, however, do not reproduce as a group (mediated by a single reproductive queen for example). Instead, they reproduce individually, and could benefit from synergies that arise from living in groups. Maynard Smith argued (Smith, 1964) that individuals that all benefit from such synergies experience group selection as opposed to individual selection (Okasha, 2005). This point of view creates a strict binary distinction between group and individual selection, which we think is too narrow to capture the actual complexity of the natural world. Individuals in a group still reproduce individually and compete against each other. The benefits from synergies contribute to their fitness as do other factors that are independent from interactions with other group members. We therefore argue that group selection applies to groups that experience synergies even if they reproduce individually, as long as the benefits that come from their interactions are restricted to the members of the group. In this way group selection is not a binary categorization but a continuum between two extrema. The remaining unknown that we address is the effect of *free riders*.

Free riders do not contribute to the success of the group, but because rewards are shared they thrive as well or better than contributing individuals. In the public goods game, for example (Davis and Holt, 1993; Ledyard, 1994), every player must make a decision about paying into a common pool. The collected amount in the pool is multiplied by a synergy factor and equally distributed among group members. In this situation, non-contributing individuals (defectors) indeed receive higher rewards than those who contributed (cooperators). In addition, they further harm the payoff to the group, since their contribution is missing during the application of synergy. Here, we will investigate what we term *pure* free riders: those players who receive the same reward as if they were a cooperating player without actually contributing to the success of the group.

The most promising solution to the free rider problem is either costly punishment (Fehr and Gächter, 2002; Boyd et al., 2003; Hintze and Adami, 2015)—allowing other

members to reduce the free rider’s payoff by paying a cost—or shunning (Panchanathan and Boyd, 2004)—excluding free riders from the group in the future. These approaches always assume that payoff is either earned individually, or distributed at the group level. In natural organisms, however, the situation is often more complicated, as resources are neither completely shared nor consumed individually. In the case of hyenas, for example, low-ranking individuals occasionally engage lions who are guarding the fresh kill of a larger animal (in our example, zebras) (Lehmann et al., 2017). The hyenas help each other to intimidate (or annoy) the lions until the lions leave their prey behind yielding it to the hyenas. In this case, the lower ranking individuals are the ones engaging the lions and thus taking the most risk (Holekamp et al., 1997). Therefore, collaboration leads to a higher group payoff group than an individual can obtain alone, but the hyenas that do not take the risk of of injury still benefit from their endeavor. At the same time, a hyena abstaining from cooperation is not necessarily getting a higher reward than those who cooperate. She might, however, avoid potential harm.

Individual Selection and Group Selection

Previous models (Traulsen and Nowak, 2006; West et al., 2007)—as far as we know—investigated group selection as a potential motive for cooperation by comparing group selection to individual selection. This strict distinction, in our opinion, does not reflect what natural organisms experience. In the hyena example from above, selection does not only happen on the group level or on the individual level alone. The zebra reward is not shared equally among all hyenas, and not all hyenas incurred identical risk. In the collaborative hunting situation, those hyenas who abstain will likely obtain less food than those hyenas that engaged the lions and thus had faster access to the zebra. These free riding hyenas then benefit less from the group hunting effort than defectors in a public goods game. Therefore, we use a proportional model where we can vary the contribution of group vs. individual selection. We use a ratio r which determines the degree to which an individual receives reward by individual actions vs. the group reward (see Equation 1). We test how virtual hyenas controlled by Markov Brains (Hintze et al., 2017) that can either engage in a collaborative hunt or just reap rewards as free riders evolve cooperative behavior given different ratios of group vs. individual selection (r).

We will show that in situations where most of the benefits for individual actions are shared with the group, then free riders thrive yet cooperation still evolves. However, when sharing only some resources, then only some free riding is allowed and evolution of cooperation is very limited.

Methods

We use Markov Brains (Edlund et al., 2011; Marstaller et al., 2013; Hintze et al., 2017) (MBs) to control the behavior of

Table 1: Computational elements from Genetic Programming.

command	description
CONST	constant value output
ADD	addition operator
SUB	subtraction operator
MUL	multiplication operator
DIV	division operator
SIN	sine operator in radians
COS	cosine operator in radians
LOG	logarithm
EXP	exponent function

the agents in the computational model. MBs are networks of computational elements which are connected to the sensors and motors of the agent in order to control it. The gates of the MB can also store and retrieve information from hidden states (memory) and their function and connectivity is encoded by a genome. Mutation, insertions (gene duplications), and deletions serve as the source of variation in this model. By using a genetic algorithm the genomes of the agents can be optimized over time, using the performance of the agents as the measure of fitness. Typically MBs use deterministic and probabilistic logic gates as their computational elements, but in this implementation we augmented the MB with mathematical operators known from genetic programming (Koza, 1994) (see Table 1) and neural computational units from NEAT (Smith, 2015) (see Table 2). These computational elements use continuous values as inputs, which are discretized for the deterministic and probabilistic logic gates such that values less than or equal to 0 become 0, and all others are considered an input of 1.

We used populations of 100 agents—referred to as hyenas—that reproduced asexually over 20,000 generations. Each generation is evaluated four times, and for each evaluation the population is randomly sorted into groups of four. Each group is then evaluated by itself. Each individual of the group is placed into the corners of a virtual environment. This environment is an area surrounded by a wall to ensure that hyenas do not randomly wander off. In the center a food resource (zebra) is placed which is surrounded by obstacles (lions). This layout is shown in Figure 1. Hyenas can sense their environment through forward facing sensors (see Figure 2) and can move in the environment by stepping forward, or by turning left or right. Hyenas can obtain resources by colliding with the resource placed in the center. However, this resource is initially surrounded by obstacles (lions) which must be pushed away. A hyena can not push an obstacle alone, however. Instead, two hyenas have to push at the same time to move the obstacle. The size of the area can vary; here we use three different sizes: 9×9 , 11×11 , and 13×13 . Observe that the smaller the area is, the easier the

Table 2: Computational elements from NEAT; the inputs become aggregated into a .

command	description
CONST	constant value output
ABS	absolute operator
CLAMP	clamps a so that $-1.0 < a < 1.0$
CUBE	a^3
EXP	$\exp a$
GAUSS	$\exp(-5.0a^2)$
HAT	hat function of a
EQU	a
INV	$\frac{1}{a}$
LOG	$\log a$
RELU	0 for $a < 0.0$ otherwise a
SIG	$\frac{1.0}{1.0 + \exp(-a)}$
SQRT	\sqrt{a}
TANH	$\tanh a$
SOFTPLUS	$0.2 \log 1.0 + \exp a$

food source can be found, and interactions between hyenas are more likely to happen by chance.

The hyenas can see what is in front, front left, and front right of them (empty, wall, lion, zebra, or hyena), and as actions they can either turn left, turn right, move forward (which is also interpreted as a push or an attempt to eat), or do nothing (see Figure 2).

The hyenas then have 200 turns in their environment, during which they can try to obtain as much of the resource as possible. Remember the resource must first be made accessible by using cooperation to move the lions blocking access out of the way. It is important to note that pushing a lion out of the way only usually only leaves one space empty from which a hyena could eat the zebra. However, two hyenas must be pushing at the same time in order to move a lion. This means that while this is a cooperative task, only one hyena at a time will be able to reach the zebra after a lion is pushed (see Figure 3).

After 200 turns, the trial ends. Over all 200 moves the number of times a hyena collides with the zebra, which is interpreted as eating, is tallied. In the case of individual selection ($r = 0.0$) this score (S_i) would now be averaged for each hyena over all four replicate experiments. In the case of group selection ($r = 1.0$) the average number of eating events ($A = \frac{1}{4} \sum_{i=1}^4 S_i$) that were observed for the entire group is given as the score for each group member. This would allow non-participating individuals to obtain the same score as those who did.

Depending on r (the ratio between individual and group payoff), we can now compute the payoff each individual receives:

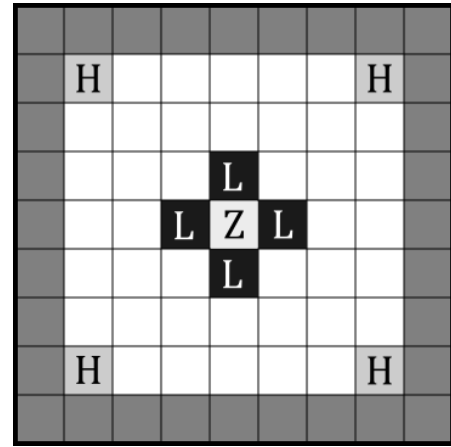


Figure 1: The starting conditions of the simulation environment. The area in which the hyenas can roam is surrounded by an impenetrable wall. Hyenas start in each corner (represented by an H). In the middle the hyenas can find food in the form of a zebra (represented by a Z). However, the zebra is surrounded by lions (represented by an L). Lions can be moved out of the way if two hyenas cooperate and push at the same time.

$$W_i = (1 - r)S_i + rA \quad (1)$$

This allows us to select the degree of group and individual selection the population experiences.

Based on this payoff computed for each individual in the population, we select the next generation proportional to their performance (using a roulette wheel selection, also known as a Moran process (Moran, 1958)). Each organism that goes into the next generation is experiencing point mutations (0.005 per site mutation rate) and deletions or gene duplications (0.00002 per site chance for either event). Observe that these mutations are experienced by the genome of each agent, which then gets translated into the brain of each agent at the beginning of each new generation. Experiments were run long enough so that we have 20,000 generations between the first generation and the most recent common ancestor. At the end of each evolutionary experiment, the line of descent (LOD) is reconstructed (Lenski et al., 2003). The LOD represents the sequence of inheritance from the most recent common ancestor back to the first generation.

Results

We first verified that after 20,000 generations, populations converged on a solution. We find that after about 15,000 generations no fitness improvements are observed (data not shown).

When evolving populations for different degrees of r (ranging from 0.0 to 1.0 in increments of 0.1) we find that

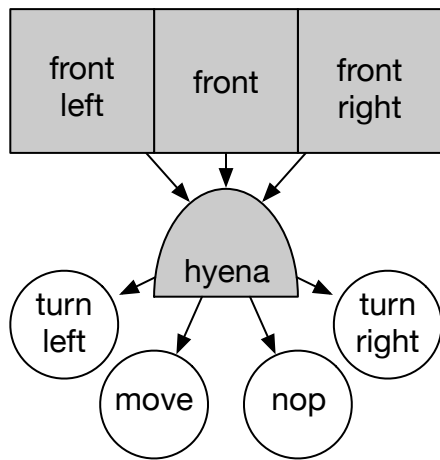


Figure 2: Hyenas can sense their environment through three sensors. Since the hyenas are positioned on a grid, they can sense what exist in the cells in the front, front left, and front right of them. They detect if a cell is empty, a wall, a zebra, another hyena, or a lion. After a hyena has received all information, it can then decide to do nothing (No Operation, or NOOP), or to turn left or right, or to move forward.

neither individual selection ($r = 0.0$) nor group selection ($r = 1.0$) results in optimal performance (see Figure 4). At the same time, in all experiments $r = 0.3$ produced the worst results. However, the r value at which performance peaks moves lower and lower with the size of the environment. For the smallest area (size= 9×9) we see that the performance peaks at $r = 0.9$. The larger the environment becomes, the lower the total score over the range of r , and the peak performance is achieved at lower r . Obviously, the smaller the environment is, the easier it is for the hyenas to find the lions and the zebra, resulting in a higher total payoff. This implies that with higher total payoffs, more can be shared with the free riders. Inversely, if it is already hard to obtain food (large environments) sharing that food becomes more problematic, resulting in a lower value for r under which optimal performance evolves.

However, these results do not imply that cooperation is at its maximum with $r = 0.6$ for large areas and $r = 0.9$ for small areas, but only that average performance is optimal. It is still possible that at the peak performance only two individuals do all the work while the two others still free ride. Therefore, we also measured the total number of hyenas in every group that attempted to push a lion at least once (even if the push is not successful). We find that generally the number of hyenas who actively seek to push steadily increases with r regardless of the size of the environment (see Figure 5). At the same time, we observe that in the largest environment (size = 13×13), we must have $r > 0.5$ before the third hyena engages, similar to a step function. Interest-




Figure 3: Schematic view of lions being pushed away from zebras. Panel 1: Lions (orange box) are blocking the access to the zebra (black striped box) for the hyenas (gray cones). Panel 2: after both hyenas pushed, the lion is moved out of the way. Now, because of random execution order, either of the two options, panel A or panel B, will happen. In either case, only one hyena will get access to the zebra. Observe that for panel A the hyena on the right needed to step forward once. In panel B the hyena coming from below stepped forward and then turned left.




Figure 4: The average fitness of population members after being evolved for 20,000 generations. Three different area sizes were used: 9×9 (solid line), 11×11 (dashed line), and 13×13 (dotted line). Gray shadow represents the standard error. Averages were made over 300 independent evolutionary experiments.

ingly, the average number of pushing hyenas stays below 3.5, indicating that not all hyenas cooperate, and between one and two continue to free ride. Apparently, the one free riding hyena does not prevent cooperation from evolving, even in complete group level selection ($r = 1.0$).

Counting the mean number of hyenas that end up eating from the zebra further supports the idea that free riders are not very harmful in this situation. We find that even for high r and regardless of environment size, only one hyena on average ends up eating from the zebra, and not all four (see Figure 6).




Figure 5: The average number of hyenas in a trial who attempted to push a lion at least once, after 20,000 generations. Gray shadow represents the standard error. Averages were made over 300 independent evolutionary experiments.




Figure 6: The average number of hyenas in a trial who actively eat from the zebra, after 20,000 generations. Gray shadow represents the standard error. Averages were made over 300 independent evolutionary experiments.

Discussion and Conclusion

We introduced the idea that selection does not only need to work on the level of the group or individual, but that selection can be a mix of individual and group contribution. This allowed us to explore the concept that group selection is one way to overcome the tragedy of the commons and to evolve cooperation. We find that neither individual nor group selection alone is optimal, but rather that sharing a large part of the payoff while still rewarding the cooperators more than the free riders is ideal.

This also intuitively makes sense. In the case of pure indi-

vidual selection, cooperation does not evolve because helping another hyena does not always pay off, since this other hyena could be the only one who gets to eat. Similarly, in the case of pure group selection, the incentive for the cooperator is missing, and thus it should also not be the ideal situation. However, without the mechanism of sharing rewards, there is no incentive to help anyone in the first place, and that is why group level selection drives cooperation.

The puzzling observation here is that cooperation evolves best when the majority of the rewards are shared, and not if the difference between cooperator and free rider is largest. In human societies we seem to be obsessed with the idea of discouraging free riders. In public goods games humans punish free riders even at considerable costs (Fehr and Gächter, 2000; Sefton et al., 2007; Carpenter, 2007; Nussair and Tan, 2011), much higher than what game theory would predict (van Damme et al., 2014). Our human social norms (Sherif, 1936) deviate in this instance from what the evolutionary optimum suggests. Since it is not clear how these norms were formed (socially or evolutionary) we can not say if they keep us from supporting free riders for evolutionary or social reasons. However, we have to question their validity and maybe consider the idea that we should be less averse towards free riders.

In this model, only one food source was used, and free riders had no alternative of finding other resources, but also did not get a benefit from abstaining from danger. We will explore the effects of alternative food sources, as well as risk from lions, and potential punishment from other hyenas in the future. Secondly, we only showed that a mix between individual and group payoff is driving cooperation in this specific situation. We will use other tasks and test if this kind of payoff scheme has similar effects on the evolution of cooperation under different circumstances. We expect to observe the same phenomenon, but the exact optimal ratio (r) to evolve cooperation might be different.

Acknowledgements

We thank Jory Schossau for insightful comments, discussions, and editing the manuscript.

References

- Adami, C., Schossau, J., and Hintze, A. (2016). Evolutionary game theory using agent-based methods. *Physics of life reviews*, 19:1–26.
- Boyd, R., Gintis, H., Bowles, S., and Richerson, P. J. (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences*, 100(6):3531–3535.
- Carpenter, J. P. (2007). Punishing free-riders: How group size affects mutual monitoring and the provision of public goods. *Games and Economic Behavior*, 60(1):31–51.
- Davis, D. D. and Holt, C. A. (1993). *Experimental economics*. Princeton university press.

- Dawkins, R. (2016). *The extended phenotype: The long reach of the gene*. Oxford University Press.
- Edlund, J. A., Chaumont, N., Hintze, A., Koch, C., Tononi, G., and Adami, C. (2011). Integrated information increases with fitness in the evolution of animats. *PLoS Computational Biology*, 7(10):e1002236.
- Fehr, E. and Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994.
- Fehr, E. and Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415(6868):137.
- Hardin, G. (2009). The tragedy of the commons. *Journal of Natural Resources Policy Research*, 1(3):243–253.
- Hintze, A. and Adami, C. (2015). Punishment in public goods games leads to meta-stable phase transitions and hysteresis. *Physical biology*, 12(4):046005.
- Hintze, A., Edlund, J. A., Olson, R. S., Knoester, D. B., Schossau, J., Albantakis, L., Tehrani-Saleh, A., Kvam, P., Shengeman, L., Goldsby, H., Bohm, C., and Adami, C. (2017). Markov brains: A technical introduction. *arXiv preprint arXiv:1709.05601*.
- Holekamp, K. E., Smale, L., Berg, R., and Cooper, S. M. (1997). Hunting rates and hunting success in the spotted hyena (*crocuta crocuta*). *Journal of Zoology*, 242(1):1–15.
- Koza, J. R. (1994). Genetic programming as a means for programming computers by natural selection. *Statistics and computing*, 4(2):87–112.
- Ledyard, J. O. (1994). Public goods: A survey of experimental research.
- Lehmann, K. D., Montgomery, T. M., MacLachlan, S. M., Parker, J. M., Spagnuolo, O. S., VandeWetering, K. J., Bills, P. S., and Holekamp, K. E. (2017). Lions, hyenas and mobs (oh my!). *Current Zoology*, 63(3):313–322.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.
- Marsteller, L., Hintze, A., and Adami, C. (2013). The Evolution of Representation in Simple Cognitive Networks. *Neural Computation*, 25(8):2079–2107.
- Moran, P. A. P. (1958). Random processes in genetics. In *Proceedings of the Cambridge Philosophical Society*, volume 54, page 60.
- Noussair, C. N. and Tan, F. (2011). Communication and punishment in voluntary contribution experiments. *Journal of Public Economic Theory*, 13(5):661–693.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *science*, 314(5805):1560–1563.
- Okasha, S. (2005). Maynard smith on the levels of selection question. *Biology and Philosophy*, 20(5):989–1010.
- Panchanathan, K. and Boyd, R. (2004). Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature*, 432(7016):499.
- Sefton, M., Shupp, R., and Walker, J. M. (2007). The effect of rewards and sanctions in provision of public goods. *Economic inquiry*, 45(4):671–690.
- Sherif, M. (1936). The psychology of social norms.
- Smith, A. W. (2015). Neat-python. [Online; accessed 10-31-2017].
- Smith, J. M. (1964). Group selection and kin selection. *Nature*, 201(4924):1145.
- Traulsen, A. and Nowak, M. A. (2006). Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952–10955.
- van Damme, E., Binmore, K. G., Roth, A. E., Samuelson, L., Winter, E., Bolton, G. E., Ockenfels, A., Dufwenberg, M., Kirchsteiger, G., Gneezy, U., et al. (2014). How werner güth’s ultimatum game shaped our understanding of social behavior. *Journal of economic behavior & organization*, 108:292–318.
- West, S. A., Griffin, A. S., and Gardner, A. (2007). Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of evolutionary biology*, 20(2):415–432.
- Wilson, D. S. and Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and brain sciences*, 17(4):585–608.