Emergence of Cooperation in Group Interactions: Avoidance versus Restriction

The Anh Han[1], Luís Moniz Pereira [2] and Tom Lenaerts[3,4]

[1] School of Computing, Teesside University, Middlesbrough, UK
[2]NOVA LINCS, Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Portugal
[3]AI lab, Vrije Universiteit Brussel, Belgium; [4] MLG, Université Libre de Bruxelles, Belgium

Public goods, like food sharing and social health systems, may prosper when prior agreements to contribute are feasible and all participants commit to do so. Yet, free-riders may exploit such agreements (Han et al. 2013), requiring then committers to decide whether to enact the public good when others do not commit. So deciding removes all benefits from free-riders but also from those who are willing to establish the beneficial resource. Here we discuss our work published in (Han, Pereira, and Lenaerts 2015) and a journal submission in (Han, Pereira, and Lenaerts), wherein we show, within the framework of the one-shot Public Goods Game (PGG) and using methods of Evolutionary Game Theory (EGT) (Sigmund 2010), that (i) implementing extra measures, delimiting benefits to free-riders, often leads to more favorable societal outcomes, especially in larger groups and highly beneficial public goods situations, even if so doing is costlier, and (ii) when restriction mechanism is not available, participation level (i.e. how many other players commit to the PGG cooperation) plays a crucial role in the decision making of commitment proposers, for their survival as well as for promoting the emergence of cooperation. Hence, there exist ethical fine tunings to be observed whenever establishing PGGs, be they for humans or non-humans, for otherwise the supporting joint moral ground may escape from under everyone's feet.

PGG is the standard framework for studying emergence of cooperation in group interaction settings (Sigmund 2010). In a PGG, players meet in groups of a fixed size, and all players can choose whether to cooperate and contribute to the public good or to defect without contributing to it. The total contribution is multiplied by a constant factor and is then equally distributed among all. Hence, contributors always gain less than free-riders (non-contributors), leading to the destruction of cooperation, as predicted by evolutionary dynamics (Sigmund 2010). In this scenario, arranging a prior commitment or agreement is an essential ingredient to encourage cooperative behavior, as abundantly observed both in the natural world (Nesse 2001) and lab experiments (Cherry and McEvoy 2013). Prior agreements help clarify the intentions and preferences of other players (Han, Pereira, and Santos 2012; Han et al. 2015). Hence, refusing to establish an agreement may be considered as intending or preferring not to cooperate (non-committers).

In our work (Han, Pereira, and Lenaerts 2015), we extend the PGG to examine commitment-based strategies in group interactions. Namely, prior to playing the PGG, commitment proposing players ask their co-players to commit to contribute to the PGG, paying a personal proposer's cost to establish that agreement. If all the requested co-players accept the commitment, the proposers assume that everyone will contribute. Those individuals that commit, yet later do not contribute, must compensate the proposers at a cost (Han et al. 2013; Martinez-Vaquero et al. 2015).

As commitment proposers may encounter non-committers, they require strategies that can deal with this type of individuals. The simplest strategy is to not participate in the creation of the common good. Yet, this avoidance strategy, AVOID, also removes the benefits for those that wished to establish the public good. Alternatively, one can establish boundaries on the common good so that only those that have committed to make it work have (better) access or that the benefit non-contributors can acquire is reduced. This strategy is referred to as RESTRICT. We compare, both analytically and using numerical simulations, these two commitment-based strategies when facing various types of free-riders, from those who do not commit, to those who commit but later defect, to those who commit and contribute only when not having to share the cost of commitment arrangement. Our results lead to two main conclusions: (i) both strategies can promote the emergence of cooperation in the one-shot PGG whenever the cost of arranging commitment is justified with respect to the benefit of cooperation, thus generalizing results from pairwise interactions (Han et al. 2013); (ii) RESTRICT, rather than AVOID, leads to more favorable societal outcomes, in terms of contribution level, especially when the group size and/or the benefit of the PGG increase, even if the cost of restricting becomes quite excessive.

In a parallel modelling work of commitment-based strategic behaviour in the context of the PGG (Han, Pereira, and Lenaerts), we consider a different set of strategies, envisaging that a restriction measure may not always be possible as it is costly and takes additional effort to implement. Namely, before engaging in a group venture individuals often secure prior commitments from other members of the group, and based on the level of participation (i.e. how many group

members commit) they can then decide whether it is worthwhile joining the group effort (Nesse 2001; Sterelny 2012; Barrett 2003). Many group ventures can be launched only when the majority of the participants commit to contribute to a common good (Cherry and McEvoy 2013). A cooperative hunting effort (both in animals, such as lions and some birds, and in humans) usually requires a sufficient number of participants "on board" to embark (Stander 1992; Alvard and Nolin 2002). While some international agreements require ratification by all parties before entering into force, most (especially global treaties) require a minimum less than the total number of negotiating countries (Cherry and McEvoy 2013). In group or coalition formation in multi-agent systems, a sufficient number of participants needs to agree on the terms of the agreement for it to be binding (Ray 2007). In general, it appears that the required participation level depends on the nature of the problem in place. We show, again using EGT modelling, that arranging prior commitments while imposing a minimal participation when interacting in groups induces agents to behave cooperatively. Our analytical and numerical results show that if the cost of arranging the commitment is sufficiently small compared to the cost of cooperation, commitment arranging behavior is frequent, leading to a high level of cooperation in the population. Moreover, an optimal participation level emerges depending both on the dilemma at stake and on the cost of arranging the commitment. Namely, the harsher the common good dilemma is, and the costlier it becomes to arrange the commitment, the more participants should explicitly commit to the agreement to ensure the success of the joint venture. Furthermore, considering that commitment deals may last for more than one encounter, we evince that longer-lasting commitments require a greater strictness upon fake committers than short ones.

The results we obtain are in close accordance with experimental economic outcomes obtained by others, see e.g. (Cherry and McEvoy 2013; Chen and Komorita 1994). But the present work further reveals that, whenever the compensation that needs to be paid by fake committers reaches a certain threshold, increasing it does not lead to improvement in terms of cooperation levels. It implies that, when designing norms, whether in real life or a self-organizing MAS, it is not necessary to have an infinitely large compensation or sanction against law breakers, for a sufficient one is enough for a wide range of situations. Moreover, the current paper suggests the need for further behavioural experiments to explore the effects of varying the essential parameters that drive commitments.

As commitments have been widely studied in AI and Computer Science, e.g. to ensure cooperation in self-organized and distributed (large) multi-agent systems, our results provided important insights into the design of such systems whenever dealing with group interactions (Bonabeau, Dorigo, and Theraulaz 1999). For instance, the key to using the potential of self-organized multi-robot systems (Gerkey and Matari 2002) is that the robots need to ensure a high level of cooperation amongst themselves, as they may have different skill sets, in order to achieve their tasks successfully. Our group commitment results appear to provide an appropriate method to ensure cooperation: the robots can arrange commitments to ensure that a beneficial coalition of skills is obtained and the task is fairly distributed (Sarker et al. 2014). Non-committers may be ostracized from the group and the mission might not be launched if the number of committers is too low.

Summing up, ethical fine tunings must be observed when establishing the norms for Public Good Games, whether for humans or non-humans, for otherwise one risks that the inherently desired common moral ground may become unfeasible.

Acknowledgments

References

Alvard, M. S., and Nolin, D. A. 2002. Rousseau?s whale hunt? *Current Anthropology* 43(4):533–559.

Barrett, S. 2003. *Environment and Statecraft: The Strategy of Environmental Treaty-Making: The Strategy of Environmental Treaty-Making.* Oxford University Press.

Bonabeau, E.; Dorigo, M.; and Theraulaz, G. 1999. *Swarm Intelligence: From Natural to Artificial Systems.* Oxford University Press, USA.

Chen, X.-P., and Komorita, S. S. 1994. The effects of communication and commitment in a public goods social dilemma. *Organizational Behavior and Human Decision Processes* 60(3):367–386.

Cherry, T. L., and McEvoy, D. M. 2013. Enforcing compliance with environmental agreements in the absence of strong institutions: An experimental analysis. *Environmental and Resource Economics* 54(1):63–77.

Gerkey, B. P., and Matari, M. J. 2002. Sold!: Auction methods for multirobot coordination. *Robotics and Automation, IEEE Transactions on* 18(5):758–768.

Han, T. A.; Pereira, L. M.; Santos, F. C.; and Lenaerts, T. 2013. Good agreements make good friends. *Scientific reports* 3(2695).

Han, T. A.; Santos, F. C.; Lenaerts, T.; and Pereira, L. M. 2015. Synergy between intention recognition and commitments in cooperation dilemmas. *Scientific reports* 5(9312).

Han, T. A.; Pereira, L. M.; and Lenaerts, T. Evolution of Commitment and Level of Participation in Public Goods Games. *Under Review at Autonomous Agents and Multi-Agent Systems journal.*

Han, T. A.; Pereira, L. M.; and Lenaerts, T. 2015. Avoiding or Restricting Defectors in Public Goods Games? *Journal of the Royal Society Interface* 12(103):20141203.

Han, T. A.; Pereira, L. M.; and Santos, F. C. 2012. Corpus-based intention recognition in cooperation dilemmas. *Artificial Life journal* 18(4):365–383.

Martinez-Vaquero, L. A.; Han, T. A.; Pereira, L. M.; and Lenaerts, T. 2015. Apology and forgiveness evolve to resolve failures in cooperative agreements. *Scientific reports* 5(10639).

Nesse, R. M. 2001. *Evolution and the capacity for commitment*. Russell Sage Foundation series on trust. Russell Sage.

Ray, D. 2007. *A game-theoretic perspective on coalition formation*. Oxford University Press.

Sarker, M. O. F.; Dahl, T. S.; Arcaute, E.; and Christensen, K. 2014. Local interactions over global broadcasts for improved task allocation in self-organized multi-robot systems. *Robotics and Autonomous Systems* 62(10):1453–1462.

Sigmund, K. 2010. *The Calculus of Selfishness*. Princeton University Press.

Stander, P. E. 1992. Cooperative hunting in lions: the role of the individual. *Behavioral ecology and sociobiology* 29(6):445–454.

Sterelny, K. 2012. *The evolved apprentice*. MIT Press.