

The Last One Standing? - Recent Findings on the Feasibility of Indirect Reciprocity under Private Assessment

Marcus Krellner^{1,2} and The Anh Han¹
¹Teesside University, UK, ²TU Dresden, Germany
E-Mail: Krellner.Marcus@gmx.de, T.Han@tees.ac.uk

Abstract

Indirect reciprocity (IR) is an important mechanism for promoting cooperation among self-interested agents. Simplified, it means: “you help me, therefore somebody else will help you” (in contrast to direct reciprocity: “you help me; therefore I will help you”). IR can be achieved via reputation and norms. However, it was often argued that IR only works if reputations are public and does not do so under private assessment (PriA). Yet, recent papers suggest that IR under PriA is feasible, and that it has more variety and ways to improve, than have been considered before.

Indirect reciprocity is usually modeled using self-interested agents playing the donation game (Sigmund, 2016): a random agent (donor) is selected to pay a personal cost c to grant benefit b ($b > c$) to another random agent (recipient). Hence, for the population as a whole, it is best if every agent decides to pay the cost (i.e. cooperate). However, the individually preferred choice for each agent is to avoid the cost (i.e. defect), making the game a social dilemma. To maintain cooperation and prevent defection, agents may use strategies that are based on reputations and norms (Nowak and Sigmund, 1998a; Ohtsuki and Iwasa, 2006).

The reputation of the recipient determines how the potential donor should act (pay the cost or not). Norms determine what reputations the acting agent will earn. A simple version of such a strategy may state: If the reputation of the recipient is good, then donate, otherwise defect; and: if an agent donates, he earns a good reputation, otherwise a bad one. This strategy was named “scoring” (Nowak and Sigmund, 1998b). It was shown that such simple strategies, whose norms only consider actions, cannot maintain stable cooperation. With these norms, agents will earn a bad reputation if they do not donate to those with a bad reputation. This is known as the problem of justified punishment (Panchanathan and Boyd, 2003).

This problem can be solved using norms that also consider the current reputation of the recipient (so-called 2nd-order norms) or even the reputation of the donor itself (3rd-order norms). For any strategy to work, defection against bad individuals should not worsen a player’s reputation.

This holds for the so-called “leading-eight” strategies (Ohtsuki and Iwasa, 2006), which were discovered by a pioneering exhaustive analytical search (Ohtsuki and Iwasa, 2004). They include L1 “standing” (Leimar and Hammerstein, 2001) and L7 “staying” (Sasaki et al., 2017).

Yet, most previous studies on the subject assumed a simplified condition: public assessment (for an excellent review see Okada (2020a)). It means, that every agent has a single reputation value (good or bad), which is agreed upon by all. Imagine a scenario where a single agent observes the interaction and then shares its perception with all others. It may commit an error, e.g. perceiving a cooperation as a defection, causing an undeserved bad reputation. But, all still agree on a single reputation value. If however, at least two agents observe the interaction and may commit perception errors independently, they may disagree on the donor’s reputation afterwards. Thus, instead of a single public reputation, there are private *opinions*. This causes severe problems to the reputation based system since the current opinions influence the assessment of future interactions (problem of PriA). Two agents with different opinions may assess the same situation differently, even if no further errors occur (Figure 1).

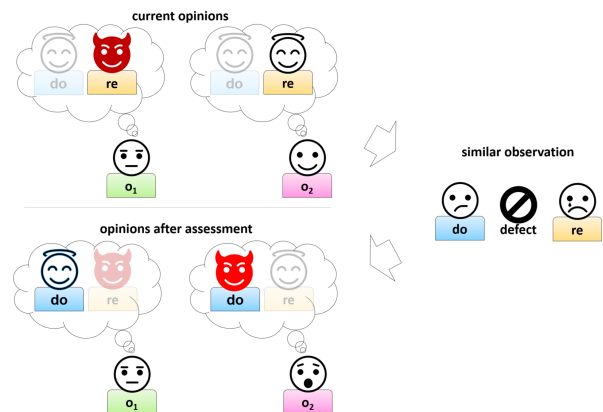


Figure 1: Disagreement about the recipient (re) between observer 1 (o_1) and observer 2 (o_2) causes new disagreement about the donor (do).

This way, bad opinions can wrongly cascade through the population, even to the point where no one will cooperate any longer. Strategies which assess cooperation against bad individuals as bad seem especially struck (Okada et al., 2017; Brandt and Sigmund, 2004), but this is a problem for all leading-eight strategies (Hilbe et al., 2018).

We will discuss recent developments in three directions. First, two new alterations of indirect reciprocity strategies have been introduced. The first of which was "pleasing" (Krellner and Han, 2020, 2021). A pleasing agent acts in accordance with others' expectations of its behavior (i.e. it pleases them), instead of being guided by its own assessment. As such, a pleasing agent can achieve a better reputation than previously considered strategies when there is disagreement in the population. It was shown that most leading-eight strategies can overcome the problem of PriA by applying pleasing. Especially L1, L6 & L7 have excellent results. Pleasing is effective even if the opinions of only a few other individuals are considered and when it bears additional costs. Pleasing is a selfish strategy, but since it increases the player's pay-off and also leads to more cooperation in the population, it causes no further dilemmas.

Another newly introduced alteration were zero-determinant strategies of IR (Schmid et al., 2021a). Such strategies can assess actions probabilistically and can therefore show some lenience. For example, instead of assessing every defection as bad, they assess 10% of them as good. This could help dealing with unintentional errors, and indeed the principle works in the strategy of "generous scoring". It is based on "scoring" (Nowak and Sigmund, 1998b), a strategy simpler than the leading-eight but that was shown to be unstable in public assessment (Leimar and Hammerstein, 2001). The zero-determinant approach revives this idea. "Generous scoring" can prevail in noisy IR under PriA, in striking analogy how "generous tit-for-tat" can prevail in noisy direct reciprocity. Sadly, this kind of probabilistic generosity seems not to help leading-eight strategies under PriA, neither applied to assessment nor action (Schmid et al., 2021b).

The second direction of recent developments were reevaluations whether the leading-eight are not successful under PriA. Depending on the evolutionary parameters of mutation rate and selection strength, there can be substantial amount of cooperation (Krellner and Han, 2021), because leading-eight tend to coexist with unconditional cooperators (All-C) (Okada et al., 2017). The more All-C are in the population, the easier unconditional defectors (All-D) can spread. However, the more All-D there are, the more some leading-eight can outperform All-C (and spread since they outperform All-D anyway). This feedback causes a stable state of cooperation. Although this is not true for all leading-eight, considerable cooperation occurs for L1 "standing", L2 & L7 "staying", and somewhat for L3 & L4.

In the last direction, two investigations were able to ex-

haustively search for successful strategies, as was done for strategies under public assessment almost two decades ago (Ohtsuki and Iwasa, 2004). Conveniently, they differ in their approach. The first (Okada, 2020b) generalized previous work on replicator dynamics in populations with three strategies: the IR strategy of interest together with All-C and All-D. They found that only two groups of strategies were successful. First, a group of four, the leading-eight strategies: L1, L7, L3 & L4. Second, a group of eight, novel strategies, which all ignore defections against cooperators. This was a surprising result, and the authors dubbed it a second way to overcome the problem of justified punishment (by not letting it be an issue in the first place). However, L1 & L7 seem strictly better than any of these new strategies.

The second paper conducted an exhaustive search for evolutionarily stable strategies (ESS) (Perret et al., 2021). This approach pitches all strategies together, considering the states of single mutants in populations of infinite residents. There was no cooperative ESS in the presence of errors. This may not surprise, because being ESS, i.e. having at least the same fitness as 257 possible invading strategies, is much harder than competing against just All-C and All-D. Especially, since the state of stable coexistence with All-C cannot be considered. If errors were excluded however, three groups of cooperative strategies were found to be ESS. The most successful group contained only L1, L2, L3 & L4. The other strategies were novel, with no overlap to the strategies found for replicator dynamics (Okada, 2020b).

The implications of both exhaustive searches are limited. They rely on deterministic evolutionary dynamics and the assumption of solitary observations (Okada et al., 2018). Yet, it is striking how some leading-eight emerge again in both of them. And, there is also overlap with results of probabilistic dynamics (Hilbe et al., 2018; Krellner and Han, 2021). Some strategies reappear in several studies: L2, L3, L4 & L7. Only L1 "standing" did so in all four (although, since no strategy is ESS in the presence of errors, L7 "staying" not being ESS should be taken with a grain of salt). L1 was actually the first strategy ever proposed to solve the problem of justified punishment (Leimar and Hammerstein, 2001). And, it should be considered a 2nd-order strategy (and so should L4, L7 & L8), because it does not use 3rd-order information (the donor's reputation). Instead, it simply ignores defections against bad recipients by all donors (the same as "keep" in Okada (2020b)). This does not require an opinion about the donor, it does not even require knowing the donor's *identity*. So, is L1 the last one ... standing?

Given the current findings, L1 but also L2, L3, L4 and especially L7 should be considered as potentially capable of indirect reciprocity under private assessment. And there might be more strategies, which were formerly not discovered under public assessment. In addition, new approaches such as *pleasing* and *generous scoring* give even more credentials to the validity of IR under PriA.

References

- Brandt, H. and Sigmund, K. (2004). The logic of reprobation: Assessment and action rules for indirect reciprocity. *Journal of Theoretical Biology*, 231(4):475–486.
- Hilbe, C., Schmid, L., Tkadlec, J., Chatterjee, K., and Nowak, M. A. (2018). Indirect reciprocity with private, noisy, and incomplete information. *Proceedings of the National Academy of Sciences*, page 201810565.
- Krellner, M. and Han, T. A. (2020). Putting oneself in everybody's shoes - Pleasing enables indirect reciprocity under private assessments. In *The 2020 Conference on Artificial Life*, pages 402–410, Cambridge, MA. MIT Press.
- Krellner, M. and Han, T. A. (2021). Pleasing Enhances Indirect Reciprocity-Based Cooperation Under Private Assessment. *Artificial Life*, pages 1–31.
- Leimar, O. and Hammerstein, P. (2001). Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 268(1468):745–753.
- Nowak, M. A. and Sigmund, K. (1998a). Evolution of indirect reciprocity by image scoring. *Nature* 1998 393:6685, 393(6685):573–577.
- Nowak, M. A. and Sigmund, K. (1998b). The Dynamics of Indirect Reciprocity. *Journal of Theoretical Biology*, 194:561–574.
- Ohtsuki, H. and Iwasa, Y. (2004). How should we define goodness? - Reputation dynamics in indirect reciprocity. *Journal of Theoretical Biology*, 231(1):107–120.
- Ohtsuki, H. and Iwasa, Y. (2006). The leading eight: Social norms that can maintain cooperation by indirect reciprocity. *Journal of Theoretical Biology*, 239(4):435–444.
- Okada, I. (2020a). A Review of Theoretical Studies on Indirect Reciprocity. *Games*, 11(3):27.
- Okada, I. (2020b). Two ways to overcome the three social dilemmas of indirect reciprocity. *Scientific Reports*, 10(1).
- Okada, I., Sasaki, T., and Nakai, Y. (2017). Tolerant indirect reciprocity can boost social welfare through solidarity with unconditional cooperators in private monitoring. *Scientific Reports*, 7(1):9737.
- Okada, I., Sasaki, T., and Nakai, Y. (2018). A solution for private assessment in indirect reciprocity using solitary observation. *Journal of Theoretical Biology*, 455:7–15.
- Panchanathan, K. and Boyd, R. (2003). A tale of two defectors: The importance of standing for evolution of indirect reciprocity. *Journal of Theoretical Biology*, 224(1):115–126.
- Perret, C., Krellner, M., and Han, T. A. (2021). The evolution of moral rules in a model of indirect reciprocity with private assessment. *Scientific Reports*, 11(1):1–10.
- Sasaki, T., Okada, I., and Nakai, Y. (2017). The evolution of conditional moral assessment in indirect reciprocity. *Scientific Reports*, 7(1):41870.
- Schmid, L., Chatterjee, K., Hilbe, C., and Nowak, M. A. (2021a). A unified framework of direct and indirect reciprocity. *Nature Human Behaviour*.
- Schmid, L., Shati, P., Hilbe, C., and Chatterjee, K. (2021b). The evolution of indirect reciprocity under action and assessment generosity. *Scientific Reports*, pages 1–14.
- Sigmund, K. (2016). *The calculus of selflessness*. Princeton University Press.