Why Might Isn't Right

Anders Sandberg
asa@nada.kth.se

December 2, 1998

Abstract

Is there an advantage in cooperating with others, even if they are weaker than you, compared to simply forcing them to do your bidding? A simple model is explored, and the results suggest that cooperative strategies that make profit independent of individual coercive strength will outdo "might is right".

The iterated prisoners' dilemma has been widely studied as a paradigm for the evolution of cooperative behavior among non-altruist agents. Under the influence of evolution reciprocal cooperation becomes dominant, eventually outcompeting mutual defection [1, 2].

In the prisoner's dilemma the agents are assumed to be equally matched; each agent is able to choose between defection and cooperation freely. But what happens if some agents are stronger than other agents, and can force weaker agents to cooperate with them while they defect?

In this paper I study a simpler model than the prisoner's dilemma, where a coercive defection strategy competes with a non-coercive cooperative strategy. Individuals switch strategy when they find their current strategy lacking.

In cooperation studies many different factors favoring and stabilizing cooperation has been found, such as a long timeframe, communication [3], reputations [4], preferential partner selection [5], group size and structure [6], reinforcement learning [7] and increasing investments [8]. It is assumed here that the cooperative agents employ some or all of these factors internally in order to stabilize their own cooperation. But is reciprocal cooperation enough to withstand outside coercion, even when the coercive agent is stronger than the coerced agent and the cooperators cannot act in any mutual defense?

1 Model

The model I study in this paper describes a population where agents can either cooperate with each other on their respective goals, or force others to work on their own goal. It is assumed that work requires two agents, even if just one is benefitting from it. Each agent has a coercive strength, and when two agents meet the one with the highest strength is able to force the other to do whatever it wants.

The model consists of a population of agents, each having a strength value, a score and a strategy. The goal is to maximize the score, which is increased by doing work or having done work for others. At each iteration, random pairs of agents are selected. One, the initiator, has the choice of acting according to its strategy. The two strategies studied are:

Might Is Right (MIR) If the other agent is weaker than the initiator, force it to do one unit of work for the initiator. Otherwise, do nothing.

Cooperate With Cooperators (CWC) If the other agent follows CWC, cooperate with it so that both agents gain half a unit of work.

Note that being forced to work for others does not imply a cost in this basic model, just a lost opportunity. In an extension of the model discussed below, a cost of being a victim is added.

Agents switch strategy when they do not earn any points in a given interval of time.

The average expectation value over the whole population of one encounter is $E(S_{\mathrm{MIR}}) = 0.5$ for MIR and $E(S_{\mathrm{CWC}}) = p$ for CWC, where p is the proportion of other CWC agents. This at a first glance suggests that MIR is the most profitable strategy for $p < 0.5$. However, for an agent of strength s, $E(S_{\mathrm{MIR}}(s)) = 0.5(1+\mathrm{erf}(s/\sqrt{2})$ while $E(S_{\mathrm{CWC}}(s)) = p$, so from the perspective of a comparatively weak agent CWC can be the most profitable strategy.

2 Simulation

A population of 1000 agents was initialized with randomly chosen strategies and strengths from a normal distribution with mean 0 and variance 1. For 50 or more iterations 5000 pairs were chosen and allowed to interact according to their strategies. At the end of each iteration, agents could shift strategy. Agents whose strategies had not worked (i.e. they had zero or less score) shifted strategy with 50% chance.

Further experiments were done using variants of the model:

One experiment was done where the strength of each agent was increased according to its score at the end of each iteration: strength $\leftarrow$ strength $+k*$score where k was a constant.

A second experiment introduced a cost of being the victim of a MIR encounter.

A third experiment allowed mixed strategies. Instead of playing pure CWC or MIR each agent chose CWC with a certain probability (and otherwise MIR). Cooperation was initiated with agents whose probability was higher than a threshold level of 0.5. This corresponds to a selective partner choice similar to the image based choice of [4].

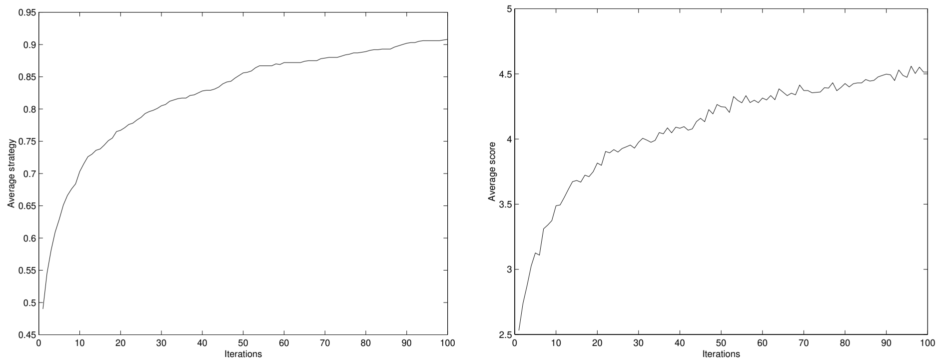

(a) Strategy

(b) Score

Figure 1: Plot of the proportion of agents following CWC and average score over time

Finally, a variant of the basic model that included CWC agents "walking into the lion's lair" was tested: if a CWC agent tried to cooperate with a MIR agent, the MIR agent would coerce the CWC agent if it was weaker.

Some alternative methods for updating the strategies (such as switching with a probability dependent on the score) were tested, but did not produce any qualitative differences unless they caused agents to shift strategy so often that the population was homogenized towards a 50-50 mix.

3 Results

Over the span of 50 iterations the proportion of CWC increased from the initial 0.5 to 0.869, and the average score rose from 2.481 to 4.269 (Figure 1). The average score and proportion of CWC were highly correlated (r=0.9981) but the correlation between individual strength and score declined from 0.4097 at iteration 1 to nearly zero at iteration 50.

The average score of MIR agents was 3.7923 with a standard deviation of 2.0027, while the score of CWC agents was 4.3402 ± 1.4403. It turned out that the CWC agents had more evenly distributed scores than MIR. The Gini coefficient was 0.196 for CWC and 0.2597 for MIR.

The average and variance of the strength of agents following MIR was 0.8323 ± 0.3916 while the average strength of CWC agents was -0.2116 ± 1.0179, i.e. a significant overlap of strengths.

When agent strengths were increased with increments proportional to their scores CWC tended to dominate, regardless of the proportionality constant. There was no significant dependence of the proportion CWC on the constant.

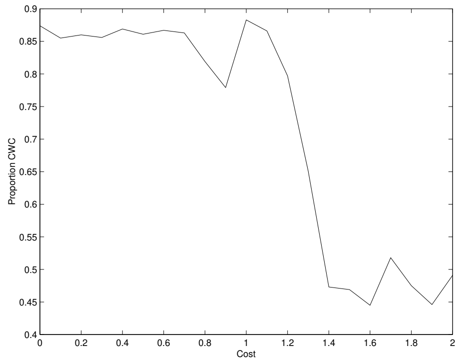

Figure 2: Proportion CWC in population after 50 iterations for different values of the penalty for being a victim to MIR.

The results of adding a cost to being the victim of a MIR coercion is shown in Figure 2. For a small cost per coercion, nothing changed substantially and CWC still won, likely because the losers quickly adopted the CWC strategy. At a cost of approximately 1 a transition took place and beyond that the population tended to homogenize towards a 50-50 mix due to strategy-jumping; agents found themselves victims and shifted strategy again and again. There was however a certain bias towards MIR, but it did not remain stable enough to produce a population dominated by the strategy.

When running mixed strategies, agents with a high probability of playing cooperatively came to dominate the population. In the end almost all agents had probabilities of playing cooperatively higher than 0.5, except for a few strong MIR-like agents (Figure 3[Mixed]). The distribution between 0.5 and 1 was roughly uniform, there was no special advantage or disadvantage in occasionally cheating by coercion. When the threshold level of cooperation was increased beyond 0.5, the population also converged to the interval above this level (not shown).

In the case of CWC initiators also being victims of MIR coercion convergence was slower than in the standard model, but CWC came to dominate the population eventually (Figure 3[Coercion]).

4 Discussion

In this simple model, mutual voluntary cooperation out-competes coercion over time.

The result is a pure CWC world, since CWC is stable against MIR invasion. MIR can't out-compete CWC and isn't stable against invasion of a team of CWC agents. On average CWC does somewhat worse in a mixed CWC-MIR world than in a pure CWC state (and MIR somewhat better than in a pure MIR

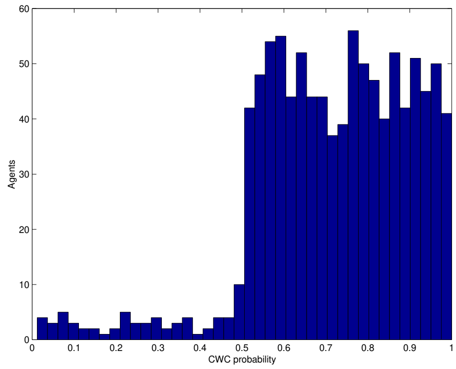

(a) Mixed

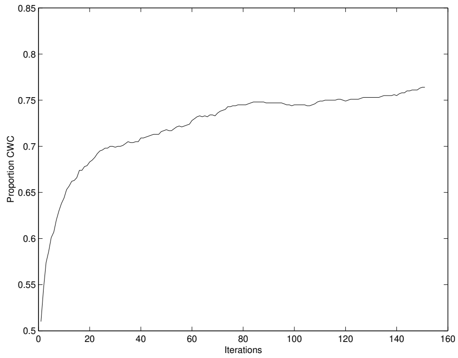

(b) Coercion

Figure 3: Mixed: histogram of the distribution of mixed agent strategies after 200 iterations. Coercion: proportion CWC in model with MIR coercion of CWC initiators.

state), but MIR doesn't do better on average than CWC in these simulations.

Strong agents are likelier to follow MIR than weak agents since on average they will benefit more. But not all strong agents follow MIR, the strength distributions overlap heavily. In general there are many strong agents that give up their strength advantage to follow CWC rather than MIR.

The effect of increasing strength depending on score not promoting MIR may at first appear counter-intuitive, since traditional wisdom says that the strong become stronger by exploiting the weak who become weaker. In this model CWC benefits because it can gain consistent large profits rather than win through strength.

The exact way agents shift strategy is important. If agents switch strategy too often, predictability is lost and the model remains a 50-50 mix of MIR and CWC.

The experiments with modifications of the basic strategy show that CWC dominance is fairly robust; it can handle both costs of being victim to MIR, the presence of MIR coercion of CWC initiators and mixed strategies. CWC is a robust strategy, related to selective cooperator strategies in the prisoners' dilemma. One reason for its success seems to be that it provides reliable profit for agents following it, regardless of their strength. This means that it is better than MIR for the weaker agents in the population, and since the value of the strategy increases with the proportion of agents following it, it will bootstrap itself into dominance, gradually reducing MIR to a minor (but rather persistent) subpopulation.

Unlike in many iterated prisoner's dilemma simulations this model assumes the agents can accurately perceive the strategies of other agents. It has been shown that reputations and visibility improves the evolution of cooperation in such cases. An interesting next step would be to extend the basic model for a wider variety of strategies to see if other results of the IPD holds in this scenario.

In this model only the non-coercive cooperators form mutually beneficial groups. Given the advantage of this, it seems likely that it would be beneficial for coercive agents to form groups too. Modifying MIR to not attack other MIR agents ironically makes it less efficient, since the number of potential victims is decreased. A MIR-CWC hybrid that uses MIR against outsiders and CWC internally can however out-compete a smaller population of pure CWC, resulting in a situation of mutual cooperation where invaders will be coerced.

It would be interesting to see what coercive strategies can form stable groups, and whether internal cheating would make them unstable. Another important question in this regard is the effect of groups that act in mutual defense, decreasing the likeliehood of successful coercion against any member.

Finally, a natural extension of this simple model would be to add coercion to the iterated prisoners' dilemma. Agents would have different strengths, and able to chose not just whether to cooperate (C) or defect (D) but also whether to attempt to coerce the other part into cooperating (producing the choices C,D,C^*,D^* where $*$ denotes coercion). The success of coercion attempts would be determined by the relative strengths of the agents.

References

[1] R. Axelrod. *The Evolution of Cooperation.* Basic Books, 1984.

[2] Robert Axelrod and Douglas Dion. The further evolution of cooperation. *Science*, 242:1385–1390, 1988.

[3] John H. Miller, Carter Butts, and David Rode. Communication and cooperation. http://www.santafe.edu/sfi/publications/Abstracts/98-04-037abs.html, 1998.

[4] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393(6685):573–7, Jun 11 1998.

[5] Dan Ashlock, Mark D. Smucker, E. Ann Stanley, and Leigh Tesfatsion. Preferential partner selection in an evolutionary study of prisoner's dilemma. Economic Report Series 35, Iowa State University, 1995. to appear in BioSystems.

[6] Natalie S. Glance and Bernardo A. Huberman. Organizational fluidity and sustainable cooperation. In K. Carley and M. Prietula, editors, *Computational Organization Theory.* Lawrence Erlbaum Associates, 1993.

[7] J. Zamora, J. R. Millan, and A. Murciano. Learning and stabilization of altruistic behaviors in multi-agent systems by reciprocity. *Biol Cybern*, 78(3):197–205, March 1998.

[8] Gilbert Roberts and Thomas N. Sherratt. Development of cooperative relationships through increasing investment. *Nature*, 394:175–179, 9 July 1998.