Color Cues and Viscosity in Iterated Prisoner’s Dilemma

Michael Joya

University of British Columbia

Abstract

The presence of color tag cues and a viscous environment have each been shown to foster conditional altruism in an iterated prisoner’s dilemma using simulated agents. The present study asks whether human subjects naturally make use of these conditions to elicit altruistic behaviour. We assume that when humans are allowed to determine both their informative cue and their cooperative strategy, no conditional altruism will emerge in either viscous or non-viscous environments. This hypothesis is disconfirmed by evidence generated by a mix of human and computer players and presented herein. Results show that both altruism and expected utility are correlated and are increased by the use of informative cues. However, this effect is attenuated by viscous environments which constrain the player to generalized strategies. It is also argued that conditional altruism is vulnerable to invasion by at least one competing strategy. Finally, it is concluded that conditional altruism requires intelligent strategy elicitation that can be directed at specific neighbours rather than applied broadly across one’s neighbourhood.

Introduction

What conditions encourage altruistic behaviour in competitive environments? Riolo, Cohen, and Axelrod (2001) have uncovered evidence via computer simulation that the presence of informative cues can establish a form of altruism in iterated prisoner's dilemma (IPD) that is not contingent on mutual reciprocity. Such evidence suggests a greenbeard effect (Dawkins 1976), which would suppose that individuals may learn to act altruistically with each other when each shares a particular genetic combination, commonly known as a greenbeard gene. We test this effect in an environment where players choose their information cue (colour) in addition to a policy specifying cooperation or defection for each of the available colours and for all opponents in a round. We hypothesize that conditional altruism will not occur in such an environment because of the mutable nature of the cue selection, as opposed to an immutable phenotype determined by genetic makeup.

Methodology

To test our hypothesis, we invited two classes of undergraduate students to provide responses for our study. Participants sat in a common area and could speak to each other, although glancing at a neighbour's screen was discouraged. Participants competed in three games each consisting of fifteen rounds. The conditions for each game corresponded with two games comprising two independent variable conditions, and a control condition with colour codes and randomized opponent pairing. The first independent variable determined the availability of colour cues for each of the players and the corresponding availability of conditional strategies for players to choose from.

The second independent variable determined whether the players repeated play against the same neighbours (hereafter referred to as viscosity) or were assigned randomized opponents each turn. At each turn, participants were asked to select a strategy specifying whether to cooperate or defect, whether to wear a red or blue colour tag, and whether to condition their actions on the opponent colour tag. In rounds 1 and 3, the use of these colour tags entailed three degrees of strategic freedom or a total of eight choices available to the participants. See figure 2. The selected strategy was applied simultaneously to each of four opponents in games that employed colour tagging, producing four separate strategy profiles per round per player. No colour tags or viscosity were present in round 2, leaving participants with only one degree of freedom (cooperate or defect) and only one strategy profile per round. No viscosity was present in games 2 and 3, meaning that each player’s opponents were chosen at random in every round.

The ordering of the rounds was reversed between the two classes of undergraduates to negate psychological biases due to recency and/or primacy. Human participants were not informed about the payoffs of the game before or during the experiment, encouraging strategy formation by trial and error. The computer participants simply played random strategies and then switched to the best-seen strategy amongst neighbours from the previous round of play.

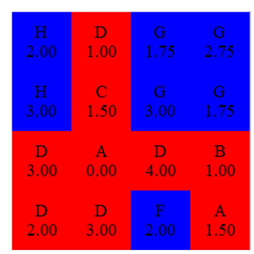

Figure 1: A Sample Round from Game 1

Each turn of play proceeded as follows. Players were asked to select a strategy during an action phase. When all players had selected their strategy, utilities were calculated for each player with respect to each of their neighbours’ selected strategies. Summary statistics are then generated for each round and displayed on a leader board showing each player's average payoff and percentage of cooperation attempts. Aliases prevent players from identifying any other player on the leader board presented at the end of each round. During this time, players were able to view the strategies of other players and connect these with their scores.

Each player's selected strategy was recorded in every turn of every round, as well as the utility gained by each player given their opponents' and neighbours' selections. Summary statistics were generated at the end of every round to describe group behaviour under the conditions presented in that round. In addition, the total and average payoffs for each player were calculated at the end of the experiment along with the percentage of games in which a player cooperated.

A summary of the experimental design is given in table 1. Note that although the round ordering was switched in the second group of participants, the round names have been kept consistent with the experimental conditions shown.

Table 1: Summary of Experimental Design

Name	Condition	Purpose
Game 1	Colour Tags, Fixed Positions	Test effect of viscosity and colour tags.
Game 2	No Tags, Random Opponents	Test effect of no colour tags.
Game 3	Colour Tags, Random Opponents	Control

To determine effects between groups, a Student's two-tailed t-test with equal variances was used to determine the differences in cooperativeness and average payoff between groups. Cooperativeness is defined as the percentage of games in which a player tries to cooperate with an opponent, and is thus a general measure of altruism for a particular individual or group. Note that this measure does not distinguish between pure altruism and conditional altruism; rather, it assesses overall attempts at cooperation whether they succeed or fail. We assume that any increases in behaviour that arise from the addition of colour tag information is attributable to a conditional form of altruism, since these conditioned strategies are what the coloured tags bring to the game.

We first compare the results of rounds 2 and 3 to test for a main effect of colour tags on the combined group of players. Note that these two rounds also differ in the number of games / opponents played per turn, in round 2 against a single opponent and in round 3 against four randomly chosen opponents. We then compare the results of rounds 1 and 3 to test for the effect of viscosity on IPD with colour tags.

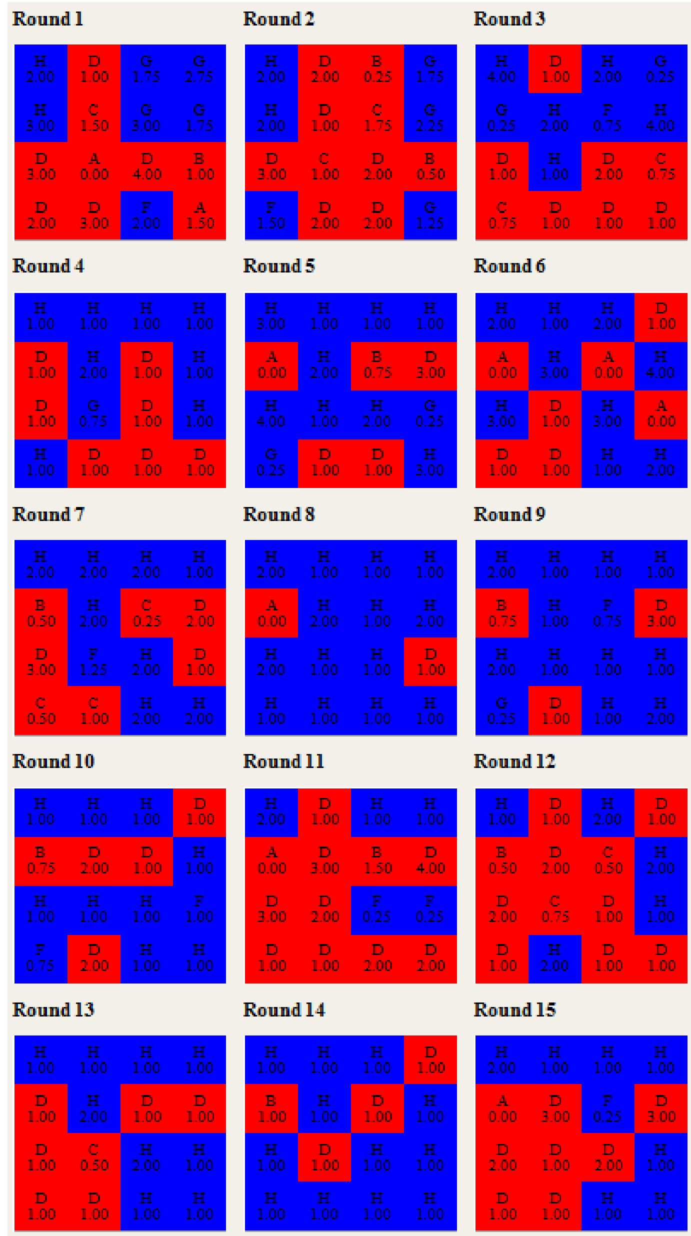

Figure 2: Game 1 (Colour tags + Viscosity)

Results

By combining the two groups of students and automatons, a subject pool of size n=32 was obtained. Each round contained 16 players, though three players / automatons were substituted and shuffled in the ISCI 422 group between games. Given that only two of these players missed two consecutive games as a result, it is not thought that this will sway our findings. The baseline results for each group are given in table 2.

Measures of cooperativeness tended to be low overall (24%) but did vary between rounds. The two participant groups did not differ significantly from one another on this measure for any of the three games played (p_1=0.32, p_2=0.44, p_3=0.10). Similarly, measures of overall average payoff reflected heavy use of defection ($1.33). These measures were only slightly above the utility received for pure mutual defection with all four neighbours ($1).

To examine whether colour tags were able to increase participant cooperativeness, we compared games 2 and 3. These games vary the availability of colour tags and hence the conditional strategies that go along with their presence. We identified a significant difference between these games for both cooperativeness and average payoff. In game 2, where colour tags were not present, participants cooperated a mere 10% of the time compared with 35% of the time when colour tags were present ($p<0.0001$). In addition, the average payoff per game was significantly higher when colour tags were present than when they were not ($1.53 vs. $1.31, $p<0.001$). These results support the hypothesis that the participants are successful at making use of colour tags to generate conditionally altruistic behaviours.

When we further augmented the environment to support altruism (both pure and conditional) by introducing viscosity, the increase in altruistic behaviour was statistically marginal. Altruism in the viscous condition was 26% compared to 35% for the non-viscous condition, which proved to be a statistically insignificant result ($p=0.180$). High variability on an interpersonal basis seemed to confound this 9% difference in altruism found between the groups. There was a mild increase in average payoff when colour tags were used in a viscous environment ($1.53 vs. $1.39, p=0.093). However, our experimental design is not equipped to answer whether this is due to the use of colour tags or simple reciprocity between neighbours, since the viscous condition was never tested without colour tags. Since conditional altruism was not suggested by the average payoffs but not by the measure of cooperativeness, skepticism remains as to the nature of the increase.

Due to the divergence between cooperativeness and average payoff seen in the last example, we decided to graph these two variables as scatter plots for each of the games to test their relation. Strangely, our control condition in game 3 seemed to be the only one in which cooperation and average payoff were not related (see figure 3.) In both game 1 (colour tags with viscosity) and game 2 (no tags, random opponents), a mild correlation exists between the tendency to defect and the average utility received per game. This tendency cannot be explained by alluding to the experimental variables in question, since both vary between game 1 and game 2. Possible reasons for this difference are discussed below, but remain speculative.

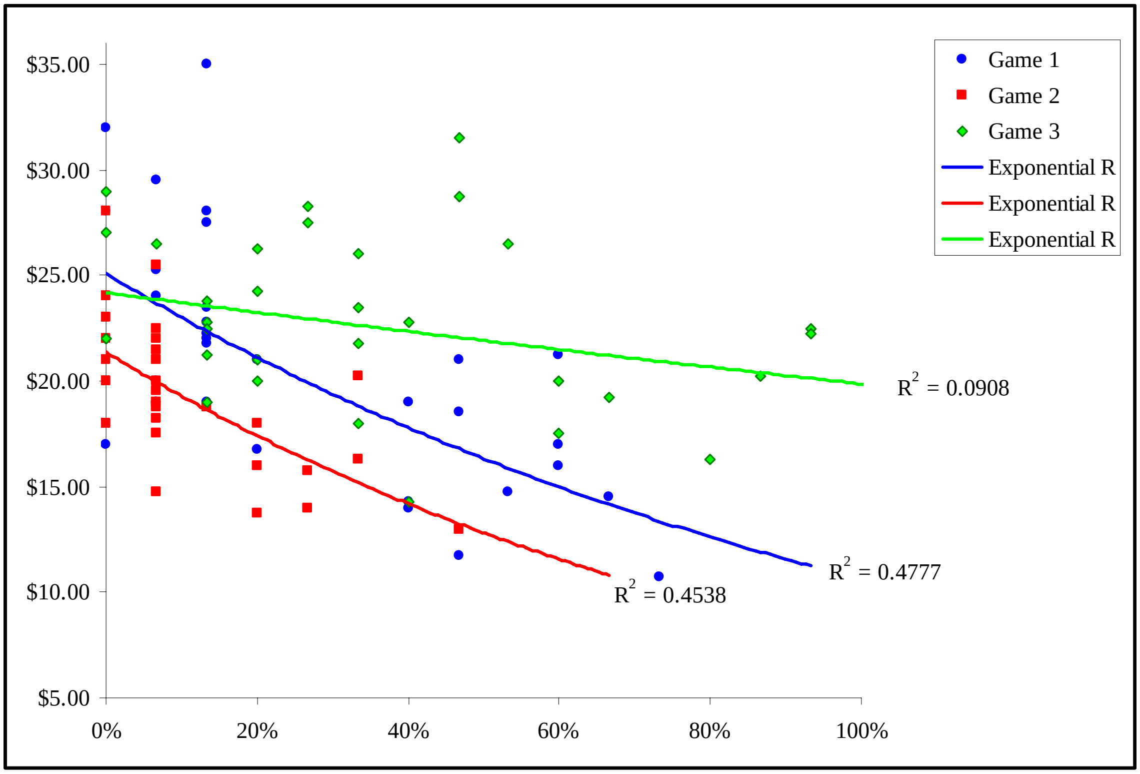

Figure 3: Comparison of Altruism and Average Payoff by Game

Discussion

The first finding among the results is that the use of colour tags does in fact foster conditional altruism between players. This is revealed by the correlated increases in attempted cooperation and in average payoff per turn when comparing games 2 and 3. It is also confirmed by the diminished slope of the regression lines when comparing game 2 with games 1 and 3 in Figure 3. The second finding is that player viscosity did not significantly raise altruistic behaviour in players using coloured tags. In fact, the correlational picture supports the view that viscosity is associated with diminished conditional altruism and average payoff for the group. This result is rather surprising in light of Axelrod, Hammond, and Grafen's study (2004) which predicts that viscosity should reinforce conditional altruism and help stabilize it via reciprocity with recently-seen neighbours.

One explanation for this result may stem from the fact that players in game 1 (the viscous condition) were not able to select specific strategies to suit past interactions with their neighbours. This made it impossible to enforce threatening policies such as tit-for-tat without harming relations with other neighbours. In a sense, this fuelled defective strategies by allowing them to become “catchy” – transferring across neighbours that were trying to punish others or protect themselves. This precipitated a situation analogous to one in standard IPD known as the “death spiral”, where two tit-for-tat strategy players can get caught in a reciprocal chain of alternating punishment. In our variant of IPD, the catalyst for this unending reciprocal punishment is fulfilled by non-neighbours who perpetuate defection in others from afar, causing a spreading wave of dissent from any (temporarily) stable islets of cooperation.

Another possible explanation for our results comes from the existence of a mixed defection strategy that can invade any localized regions of cooperation. Assume for argument that strategy X is a conditional altruism strategy that is stable in a particular region of the grid. In this region, conditional altruists recognize the value in playing either red or blue and cooperating with their own color while defecting against others. Now consider that one player in this region switches to a pure defection strategy D in order to reap rewards from like-coloured altruists in the region. Strategy D gains more utility than strategy X, but can only do so if strategy X assumes the same colour as it does. The defector can overcome this condition by switching colours 50/50. This defector strategy D’ gains higher utility against X than X does against itself, and so cannot be invaded by any of the conditional altruism strategies.

This experiment thus disproves our hypothesis that informative colour tags do not lend themselves to conditional altruism in IPD. We have shown that indeed they do foster conditional altruism, but that it is severely hampered when players unilaterally apply a strategy to all their neighbours rather than selecting strategies for neighbours on an individual basis, given past relations with that neighbour. Future study along these lines should provide players with the ability to respond to neighbours as individuals rather than as a group. In closing, we have shown that the use of informative signals such as colour tags may foster conditional altruism, but that this effect can be limited by an environment that restricts players to a unilateral strategies rather than allowing strategies to be selected on a per-neighbour basis.

Works Cited

Axelrod, R., Hammond, R., & Grafen, A. (2004) Altruism via Kin-Selection Strategies That Rely on Arbitrary Tags With Which They Coevolve. Evolution. 58(8):1833-1838.

Dawkins, R. (1976) The Selfish Gene (Oxford Univ. Press, Oxford).

Riolo, R., Cohen, M.D. & Axelrod, R. (2001) Evolution of Cooperation without Reciprocity. Letters to Nature. Vol.414:441-443.