Evolution of strategies for the collective-risk social dilemma relating to climate change

To cite this article: G. Greenwood 2011 *EPL* **95** 40006

A Letters Journal Exploring the Frontiers of Physics
August 2011
EPL, **95** (2011) 40006 www.epljournal.org
doi: 10.1209/0295-5075/95/40006

Evolution of strategies for the collective-risk social dilemma relating to climate change

G. Greenwood[(a)]

Portland State University - Portland, OR 97201, USA

received 28 March 2011; accepted in final form 5 July 2011
published online 5 August 2011

PACS 02.50.Le – Decision theory and game theory
PACS 87.23.Ge – Dynamics of social systems

Abstract – This paper describes an N-person social dilemma game created to study how climate change agreement strategies develop. The players in this game are trying to collect a target amount for a climate aid fund. Players (countries) fall into three categories: cooperators who contribute their fair share, altruists who contribute more than their fair share and defectors who contribute nothing. In all cases we would evolve a set of player strategies that collected the target sum ($-0\%/+0.5\%$). Our results indicate defection is a preferred strategy, but trigger strategies can markedly improve contributions. Our game is designed to see what motivates countries to live up to their agreements to contribute to climate change aid funds.

In December 2009 the United Nations Framework Convention on Climate Change (UNFCCC) met in Copenhagen, Denmark, to set the stage for a new climate protocol to replace the 1997 Kyoto Protocol [1], which was due to expire. Nearly 200 countries attended.

The financial burden developing countries would face as they contend with the effects of *anthropogenic global warming* (AGW) was an issue raised at the conference. Near the end of the conference a proposal was introduced asking richer countries to establish a climate change aid fund to help developing countries offset these costs. The US Secretary of State, Mrs. Clinton, announced the United States is willing to participate in a climate change aid fund of up to 100B USD by 2020 —providing certain other countries also commit to participating, although the exact US contribution was not stated. Australia, Great Britain, Norway, Japan and France also agreed to participate while other countries said they would participate only if a formal agreement was signed.

Social dilemmas arise whenever a group of individuals must decide how to share a common resource while balancing short-term self-interests against long-term group interests. Put another way, social dilemmas are situations where individuals must choose whether to cooperate for the good of the group or to defect for personal gain, bearing in mind the whole group suffers if everyone defects.

(a) E-mail: greenwd@ece.pdx.edu

AGW is sometimes portrayed as a tragedy of the commons problem [2]. Every country contributes to the global levels of greenhouse gas (GHG) emissions, but the negative effects on individual countries are negligible. Hence, countries have little incentive to voluntarily reduce their GHG emissions given that the associated costs may be detrimental to their economy.

Participation in climate change programs qualifies as a tragedy of the commons problem. Everyone is vulnerable to over-exploitation of the Earth's climate according to Pfeiffer and Nowak [3]. Cooperation buys goodwill from other countries but the negative impact to a country's economy may be too much to ignore. Countries may therefore be reluctant to voluntarily reduce their GHG emissions. But fines or other penalties for non-compliance exist even at international levels. For example Article 18 in the Kyoto Protocol [1] called for "an indicative list of consequences, taking into account the cause, type, degree and frequency of non-compliance". Hence there is pressure to cooperate by contributing after all. Countries may also feel pressured to contribute if they see other countries contributing.

Deciding whether or not to contribute to climate change aid funds is a different kind of social dilemma, more in line with a public goods problem. Some countries may contribute less than their fair share to aid funds because they hope other more altruistic countries will make up for any shortfalls. Of course the target sum could never be attained if every country made that decision.

Game theory provides a perfect framework for studying strategies countries use to decide whether to cooperate or defect in these real-world social dilemma situations. One of the best examples involving climate change social dilemmas is the recent work by Milinski and his colleagues [4]. They conducted a series of experiments relating to climate change agreements. A large number of university undergraduates were recruited to participate in a public goods game. Students were given an endowment and asked in each of 10 rounds to contribute to a "climate account". A target sum was specified and the students were told they could keep any money they did not invest. However, failure to collect the target amount could result in the loss of any uninvested money. Unfortunately, all Milinski did was conduct the experiments and then perform a statistical analysis on the results. The thought process used by the students to formulate an investment strategy was largely ignored. (Students were asked to fill out a questionnaire after the experiments, but few details were given in Milinski's papers.)

We created a social dilemma game, patterned after the Milinski experiments, to investigate how climate change agreement strategies evolve over time.

Social dilemma games have $N > 2$ players who must decide, over a series of rounds, whether to cooperate or to defect. A cooperating player receives a maximum payoff if everybody else cooperates, but a single defector can always get a larger payoff. Therein lies the dilemma: defection pays better than cooperation, but if everybody defects then the entire group suffers. Social dilemma games include N-person prisoners' dilemma and the related family of public goods games.

It is important to understand the key features of the experiments conducted by Milinski and his colleagues since our social dilemma game is based on those experiments. Only an overview of their work appears here; interested readers are referred to [4] for specific details.

The Milinski experiments were conducted to see if a group of unrelated people could reach a target sum through individual contributions. They defined a special type of social dilemma game called *collective-risk social dilemma*. Thirty groups of six undergraduate students apiece played a public goods game to simulate the collective-risk social dilemma problem. Each student got an initial endowment of €40. During each of 10 rounds a student could contribute €0, €2 or €4. All funds were put into a common "climate account". The target sum was €120 (equivalent to a contribution of €2 per player per round). Players could keep any left over funds. However, if the group failed to collect the target sum they could lose all of their savings with probability p.

The students were told the total sum of money, collected from all groups, would go towards publishing an advert on climate protection in a prominent German newspaper. All participants had previously received information on the horrors of climate change presumably to induce some desire to contribute to the climate fund. There is no indication opposing viewpoints were presented.

In general the student groups were reluctant to contribute to the climate account —even with a high probability of losing everything for failure to collect the target sum. Indeed, only 5 out of 10 groups met the target when the probability of losing everything was $p = 90\%$. (Only 1 out of 10 groups with a 50% probability and 0 out of 10 groups with a 10% probability.) Further analysis indicates students tended to contribute willingly early on, but contributions declined steadily in later rounds. Milinski *et al.* provided no explanation for this behavior.

Our game is formulated as a public goods game. Each player (country) has s strategies to choose from. Most researchers investigating multi-strategy performance find $s = 2$ is sufficient because game dynamics do not qualitatively change much for $s > 2$. A strategy tells a player how much to play (€0, €2, or €4) during each round, however his play can vary depending on how close the entire group is to meeting its target sum.

Twenty players participate in a 10 round game. Payoffs are deferred until after the 10th round. The target sum of €400 is based on a fair and equitable €2 contribution per player per round over the 10 rounds. Players who contribute €2 in a round are called "cooperators" (C) because they contributed the fair amount of €2. In other words, they are cooperating in reaching the target sum by contributing their fair share. The "defectors" (D) are those who contribute €0 and the "altruists" (A) are those who contribute €4. Since players get to keep any money left in their accounts at the end of the 10th round, defectors are free riders who intend to profit from the contributions of others. Conversely, those who contribute €4 in a round are called altruists because they contributed more than their fair share to help the group meet its target.

Each player strategy is comprised of three genes, which are interpreted as probabilities: p_1, the probability of cooperating, p_2 the probability of defecting and p_3 the probability of being altruistic. We require $\sum_i p_i = 1$. Thus each p_i is a sub-interval on the unit interval. A player's choice during a given round is made by randomly choosing a number on the unit interval and then determining which sub-interval that number resides in.

The strategies are modeled as probabilities to reflect changing economic conditions. Natural disasters and regional conflicts can have devastating effects on a country's economy. Consequently, countries may not always be able to make annual contributions to climate change aid programs despite prior commitments to do so.

At the beginning of each set of 10 rounds players receive an endowment of €40, which is deposited into a personal bank account. Each round the player makes an investment, according to the strategy used, and this investment is deducted from their accounts. At the end of the 10 rounds the investments from all players are totaled and compared against the target sum of €400. Players can keep any money in their accounts after the 10 rounds, so there is no advantage to investing more than is necessary. However, if the group fails to meet its target sum then everyone must give up 90% of the money in their account. This penalty provides an incentive to invest.

After each round a strategy receives virtual points indicating how good that particular strategy is with respect to helping the group meet its target. Strategies can receive virtual points whether or not they were actually played. Players always choose the strategy with the highest number of virtual points.

Virtual points are awarded if the strategy helps the group meet its target. If the group has acquired less than half of the target, plays that make contributions —*e.g.*, C or A— are more desirable. On the other hand if the group has already acquired more than half of the target, a player would not want to over-contribute; a strategy that plays D or C would be preferable. So, strategies that have the highest probability of playing C or A get awarded a virtual point if less than half of the target has been acquired, whereas those that play C or D with the highest probability get a virtual point if more than half has been acquired. For instance, suppose in a given round less than half of the total has been collected. A strategy that has $\max(p_1, p_2, p_3)$ equal to p_1 or p_3 would get a virtual point.

The size of a player's personal bank account is used as a measure of fitness, which will be defined shortly. After every set of 10 rounds the worst fit player is replaced by a clone of the best fit player. One of the strategies in the clone is then mutated with a 10% probability. (Mutation just randomly perturbs each gene value and then renormalizes them to ensure $\sum_i p_i = 1$.) This new player starts with a zero bank account balance, although he will get an endowment along with the other players if and when another set of 10 rounds is started.

The experimental results reported below were averaged over 100 trials. Each trial was randomly initialized —*i.e.*, the p_i gene values were randomly initialized for each player at the beginning of each trial. Since each player uses a mixed strategy, unlike some social dilemma games our game has no initial distribution of player types. Gene values do not change except every 10th round where the worst fit player is replaced with a clone of the best fit player (with possibly mutated gene values in one of the strategies). It is through this cloning/mutation process that strategies in the population evolve over time.

The fitness function must be carefully constructed. The underlying assumption is the whole group is really committed to collecting at least T euros for the climate change fund. Of course this goal could be easily achieved if players contributed the maximum amount each round, but that strategy works against a player's best self-interest. (Remember, players get to keep any left over money in their accounts if the target sum is reached.) Consequently, the fitness function should assign high fitness to players that strike a balance between their own self-interests and helping the group collect the target amount.

Let P_r represent the r-th player. Then the fitness of P_r is calculated by

$$\text{fitness}(P_r) = \begin{cases} \exp\left(\frac{\Delta k}{T}\right), & \text{total} < T, \\ \exp\left(\frac{-\Delta k}{T}\right), & \text{otherwise}, \end{cases} \tag{1}$$

where $\Delta k = (20 - \zeta)$, ζ is the player's bank account balance, "total" is the amount invested by all players after the 10 rounds and T is the target sum.

The fitness function form changes depending on whether or not the group invested at least T euros. First consider the case where total $< T$. Clearly the existing strategies do not collect enough money. We can collect more money by awarding higher fitness to altruistic strategies. €20 is an individual's fair share investment (€2 per round over 10 rounds), which explains the first term in the Δk equation. Players start each set of 10 rounds with a new €40 endowment. An altruistic strategy makes a player invest more than €22 resulting in a bank balance of $\zeta <$ €20. Hence, $\Delta k > 0$ and $\text{fitness}(P_r) > 1$, making altruistic players higher fit. Conversely, investing less than their fair share makes $\zeta >$ €20, making $\Delta k < 0$ so $\text{fitness}(P_r) < 1$. Just the opposite fitness assignments are made if the investment total reaches or exceeds the target sum because the sign of Δk is negative. Note any strategy that invests exactly the fair share of €20 gives that player a fitness of 1.0 regardless of whether the target sum was met.

It seems unlikely players will contribute exactly the correct target sum on any arbitrary 10 round play since the strategies are evolved and evolution is a stochastic process. Therefore, we relaxed the termination criteria and let the game end if the total amount collected after 10 rounds is $T \leqslant \text{total} < 1.005T$. That way the game terminates with a successful set of strategies that collects enough money to at least meet the target sum but which does not collect an excessive amount.

The first task was to define what makes a player C, D or A. An unconditional cooperative player contributes €2 in each round or €20 total over 10 rounds. Hence, we use a 10 round contribution of €20 to classify a player as a cooperator. Defecting players then contribute less than €20 and altruistic players contribute more than €20. Note the €2 per round contribution is an average. Players do not have to contribute €2 each round to be cooperators. For example, a player could contribute €0 in one round, €2 in 8 rounds and €4 in one round and still be classified as a cooperator. The game had 20 players, so if everyone cooperated the total sum collected would be €400.

As mentioned above, after 10 rounds of play the total amount contributed is compared against the target $T =$ €400. One trial ends if the termination criterion was met. Otherwise, the trial continues by clearing the total and initiating another 10 rounds of play. Players start a new 10 round set with their old strategy pairs (except for the cloned player whose strategy pairs were replaced and then possibly mutated.) Each trial, on the other

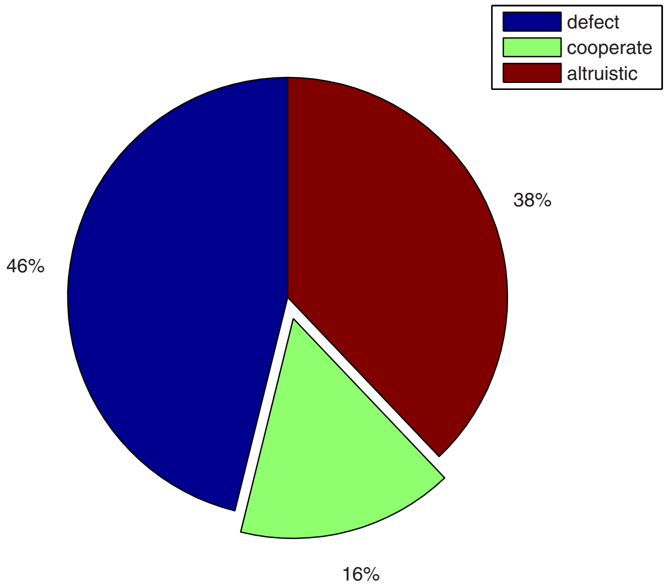

Fig. 1: (Colour on-line) Distribution of player types for the basic social dilemma game. Results are averaged over 100 trials.

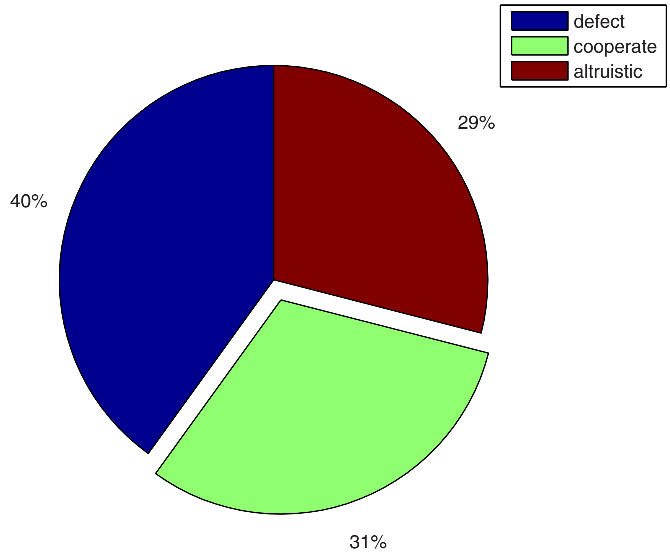

Fig. 2: (Colour on-line) Distribution of player types for the social dilemma game with trigger strategies. Four (out of twenty) players defected with a probability depending on the number of cooperators in the previous round. (See text.) Results are averaged over 100 trials.

hand, starts with every player getting a new randomly generated pair of strategies. In our experiments trials typically terminated within 250–300 rounds (25–30 sets of 10 round plays).

Figure 1 shows the results averaged over 100 trials. Cooperative strategies did evolve in roughly one out of six trials with more free riders (defectors) than altruists. Intuitively one would expect an equal number of defectors and altruists to cancel each other out. However, it is the contributions that must balance and not the number of each type of player. Altruists contribute more than €20 over 10 rounds —how much more is the only important thing. A single altruist who contributes €30 cancels out two defectors who only contribute €15 apiece.

It is useful to briefly discuss cooperation and defection on the international level before describing how our social dilemma game relates to climate change programs. Cooperation here means a country (player) agrees to deposit a specified amount into a climate change aid fund at a specified time; defection means the deposit is deferred or never made. A cooperating country derives no direct benefit from the fund and deposits are not returned. What then motivates a country to cooperate? The only obvious pay off is cooperation elicits good will from others, which can be exploited in future international negotiations.

On the other hand, countries have a strong incentive to defect because contributing to an aid fund negatively impacts their own economy. Governments do not have any money of their own. The only money a government has is collected from the private sector via taxes, fines or fees. Governments can also print money, although that leads to inflation. Either way the economy suffers. Moreover, any money contributed to a climate change aid fund does not go back into their respective economy but is spent elsewhere by other countries.

Our game dynamics resemble cultural evolution [5] in humans, where successful strategies spread via imitation. But strategy imitation in our game is neither common (only one player in the population) nor frequent (only once every 10 rounds). Moreover, the imitation is not exact because one of the cloned strategies can be mutated. These dynamics seem reasonable in the context of climate change programs.

Suppose the best fit player is inclined to defect —*i.e.*, $p_2 > p_1$ and $p_2 > p_3$. This situation exists whenever the total contribution is more than the target amount. If many other countries suddenly imitated this defect strategy, the total contribution in the next 10 rounds of play could swing well below the target amount. Now altruistic countries become more fit. Mass switching of strategies could lead to wild oscillations in the contributions, which makes it difficult for the group to collect the target amount. Conversely, if just one country imitates the defect strategy then the total contribution over the following 10 rounds of play should slightly decrease, bringing the total contribution closer to the target amount. Other countries would not have to change their strategies either, which creates certainty and therefore helps future budget planning. Sudden and unanticipated circumstances —*e.g.*, earthquakes or other natural disasters— may not permit a country to exactly imitate another strategy so a slightly altered (mutated) version might be required.

The Milinski experiments showed student groups generally did not want to contribute to a climate change fund —*i.e.*, the groups avoided cooperation. Indeed, even when the chance of losing their savings was 90%, only half of the groups collected the target sum. Cooperation evaporated at lower risk levels (50% and 10%).

The strategies that evolved for playing our social dilemma game are consistent with the Milinski experiments with human subjects. Our game shows under normal circumstances defecting strategies preferentially evolve while relatively few cooperators appear. In order to reach the target sum the population must contain a large number of altruistic players. The student groups in the Milinski experiments seldom collected the target sum primarily because few groups were cooperative, let alone altruistic.

For a second experiment we considered trigger strategies. Twenty percent of the players were designated as defectors who would not cooperate unless a minimum number (μ) of other players cooperated first. These players are unlikely to ever be altruists so p_3 was set to zero. However, if at least μ players did cooperate in the previous round then these players would fully cooperate ($p_1 = 1.0$) in future rounds. We chose $\mu = 8$, which corresponds to 40% of the population. Figure 2 shows the cooperation level doubles with trigger strategies.

These results raise serious concerns about whether countries will be long-term contributors to climate change aidbe funds. The plan is countries would make yearly contributions to the fund. However, the results from our social dilemma game suggest cooperation is not to expected and the only hope for success is if a reasonable number of altruistic countries persist. Countries are unlikely to be altruistic year after year, especially given the current state of the global economy. Consider, for example, how some members of the European Union are reluctant to help Greece, Ireland and now Portugal deal with their economic problems. Trigger strategies, however, can improve participation because pressure can be applied to countries who conditionally agree to participate if a sufficient number of other countries have already contributed.

REFERENCES

[1] Kyoto Protocol to the United Nations Framework Convention on Climate Change, United Nations (1998).
[2] Hardin G., *Science*, **162** (1968) 1243.
[3] Pfeiffer T. and Nowak M., *Nature*, **441** (2006) 583.
[4] Milinski M., Semmann D., Krambeck H. and Marotzke J., *Proc. Natl. Acad. Sci. U.S.A.*, **105** (2008) 2291.
[5] Traulsen A., Semmann D., Sommerfeld R., Krambeck H. and Milinski M., *Proc. Natl. Acad. Sci. U.S.A.*, **107** (2010) 2962.