

Stages in the Evolution of Ethnocentrism

Thomas R. Shultz (thomas.shultz@mcgill.ca)

Department of Psychology and School of Computer Science, McGill University, 1205 Penfield Avenue
Montreal, QC H3A 1B1 Canada

Max Hartshorn (max.hartshorn@mail.mcgill.ca)

Department of Psychology, McGill University, 1205 Penfield Avenue
Montreal, QC H3A 1B1 Canada

Ross A. Hammond (rhammond@brookings.edu)

Economic Studies Program, The Brookings Institution, 1775 Massachusetts Avenue, NW
Washington, DC 20036 USA

Abstract

Computer simulations suggest that agents cooperate in the Prisoner's Dilemma game, despite the individual cost of doing so, because they have evolved to cooperate, at least with individuals of the same kind. We present new simulations indicating that, early in evolution, humanitarian strategies compete successfully with ethnocentric ones. In just over half of the simulated worlds, there was an early stage in which humanitarian strategies, involving cooperation with all agent types, dominated over other strategies. Ethnocentric strategies involving only in-group cooperation eventually dominate.

Keywords: Agent-based simulation; ethnocentrism; evolution; humanitarianism.

Introduction

It has been a mystery why people cooperate so much in the infamous Prisoner's Dilemma (PD) game. Such cooperation is mysterious because the most rational course of action in a one-time PD game against an unknown opponent is to defect rather than cooperate. One of several intriguing explanations for this mystery is that humans have evolved a tendency to cooperate with members of their own ethnic group but discriminate against members of other ethnic groups, essentially an ethnocentric strategy. A computer simulation of the evolution of ethnocentrism showed that this explanation is indeed plausible, with ethnocentric genotypes stabilizing in about three-fourths of the population (Hammond & Axelrod, 2006).

Our paper uncovers another interesting aspect of this issue, namely the possibility that there can be stages in this evolutionary process. In particular, we found evidence for an early stage in which a humanitarian strategy dominates. Humanitarianism would involve cooperation with any agent, whether a member of one's own group or not.

We begin with brief reviews of the PD game and the Hammond-Axelrod simulation before presenting and discussing our simulation results on possible evolutionary stages.

The Prisoner's Dilemma Game

PD is widely considered to be an elegant embodiment of the problem of achieving mutual cooperation. It was designed in the 1950s to study global nuclear strategy with game theory. For nuclear strategists, defection would involve launching missiles towards an opponent, whereas cooperation would involve not sending them. The game was later reformulated with a story about prisoners, apparently to better interest experimental psychologists. For prisoners, defection would involve ratting out a colleague and cooperation would involve keeping quiet.

The basic characteristics of PD are illustrated in Table 1. There are two autonomous agents A and B who each have the choice of cooperating with or defecting against each other, but without knowledge of what the other agent will do. In the classic version of the game, which is studied here, these agents may interact only once so they cannot learn how the other individual agent plays.

Table 1: Characteristics of the PD game

	B Cooperate	B Defect	Mean _A
A Cooperate	$R_A = .02$ $R_B = .02$	$S_A = -.01$ $T_B = .03$.005
A Defect	$T_A = .03$ $S_B = -.01$	$P_A = .00$ $P_B = .00$.015
Mean _B	.005	.015	

If agents A and B both cooperate, they each receive a modest *reward* R. If agent A cooperates and agent B defects, A receives the *sucker's payoff* S for being the only one who cooperates, and B receives a larger reward T, the *temptation for defecting*. If agent A defects and agent B cooperates, there is a mirror-image outcome in which agent A receives T, and agent B receives S. If both agents decide to defect, then each receives P, the *punishment for mutual defection*. Table 1 has the actual values used by Hammond and Axelrod and by us, but more generally these values can vary, as long as they

obey the following inequalities: $T > R > P > S$. Deviations from this constraint may create other games, which have been much less studied than PD.

The fact that $T > -S$ derives from the assumption that the benefit of receiving cooperation (here .03) exceeds the cost of giving cooperation (here .01). This inequality is a natural consequence of social specialization in a relatively abundant environment. It costs little for a specialist to give his surplus product to another agent, who may find the gift very beneficial because he specializes in something else.

The mean outcomes for agents A and B are shown in the marginal column and row, respectively, of Table 1. The so-called *Nash equilibrium* explains that defection is the rational course of action for a player who doesn't know what the opponent will do. Notice that the mean outcome for defection in such cases is three times the mean outcome for cooperation. Thus, mutual defection is the only outcome from which each player could only do worse by unilaterally changing its move. Yet organisms from bacteria to humans are observed to largely cooperate in PD games (Axelrod & Hamilton, 1981). Hence, the mystery referred to earlier. Why do PD agents cooperate so readily, even in single-shot versions of the game with no chance for learning and reciprocity?

The Hammond-Axelrod Simulation

A computer simulation by Hammond and Axelrod (2006) provides one possible solution to the mystery, namely that PD cooperation is a byproduct of evolution, but with the interesting qualification that the resulting cooperation is not universal but rather ethnocentric. They devised an abstract agent-based simulation in which each agent possessed three heritable characteristics: a tag representing a superficial marker of membership in one of four different groups, an in-group strategy (cooperate or defect), and an out-group strategy (also cooperate or defect). Each agent also had a modifiable reproductive potential, initially a smallish proportion (.12). On every evolutionary cycle, each agent played a one-shot PD game against each of its (possibly four) neighbors. Payoffs for adjusting reproductive potential were as described in Table 1.

The Artificial World

The agents' world was a 50 by 50 lattice in which each position could be potentially occupied by a single agent. Each agent could potentially interact with each of its four neighbors – to the north, east, south, and west. This is known as the *von Neumann* scheme for creating neighbors, after John von Neumann, who pioneered agent-based simulations with his invention of cellular automata. To uniformly enforce this constraint, the lattice was effectively folded from north to south and from east to west, thus creating a donut-shaped world (technically, a torus). This conceptual folding, commonly used in such simulations, allows agents residing on the edges of the

lattice to have the same number of neighbors as all others (instead of only three neighbors, or two if residing in a corner).

Phases of an Evolutionary Cycle

Each evolutionary cycle went through four phases: immigration, interaction, reproduction, and death.

Immigration In the immigration phase, new agents were created at an immigration rate, with tag and strategies randomly assigned, and were placed at a randomly selected empty location on the lattice. Evolution began without any agents, and the default immigration rate was set to 1 immigrant per cycle.

Interaction At the start of the interaction phase, each existing agent had its reproductive potential set to the default value of .12. Then each agent A had a chance to interact with each of its neighbors N (up to four). As noted earlier, this interaction involved a one-shot PD game in which agents A and N would independently decide whether to cooperate or defect based on their own strategies and the tags of each. If agent A decided to cooperate, its reproductive potential was reduced by a cost value (default .01) and the reproductive potential of agent N was increased by a benefit value (default .03). Agent N made an analogous decision with analogous adjustments to reproductive potentials. More generally, a donor's reproductive potential lost .01, and a recipient's reproductive potential gained .03.

Reproduction The reproduction phase began by sorting the existing agents into a new random order (a standard procedure to avoid a potential source of simulation artifacts). In this order, each agent had a chance (based on its current reproductive potential) to reproduce a clone if an adjacent empty space was available. The clone, if any, was placed in a randomly-selected empty neighboring location if any existed. A clone possessed the characteristics of its parent, subject to a mutation-rate parameter (default .005) applied to each of the three inherited characteristics. Strategies could mutate to an alternate binary value (from cooperate to defect or vice versa). A tag could mutate to any one of the three other tags, randomly selected with equal probability.

Death In the death phase, each agent had a chance of dying equal to the death rate (default .10), which would result in its removal from the lattice. The default number of evolutionary cycles was 2000, chosen because stable patterns generally arise in the model well before this point.

Strategies

With two binary strategy characteristics (cooperate or defect with in-group or out-group), there were four possible genotypic strategies that could evolve. A *selfish* strategy always defects. A *traitorous* strategy cooperates with agents of other groups but not with agents of its own group. An *ethnocentric* strategy does the reverse, cooperating within its own group but not with agents of

different groups. And a *humanitarian* strategy cooperates with every kind of agent.

Results

Genotypes and phenotypes were assessed over the last 100 of the 2000 cycles. Hammond and Axelrod (2006) reported mean genotype strategy proportions of .08 selfish, .02 traitorous, .75 ethnocentric, and .15 humanitarian. In terms of phenotypic behavior, agents could either cooperate or defect, and the mean proportion of cooperation over all interactions was .74. Each parameter, including cost, number of tags, lattice size, and number of cycles was both halved and doubled in comparison runs. Genotypic and phenotypic outcomes were relatively stable, although ethnocentrism went up a bit with number of tags and down a bit with cost and mutation rate.

The simulation showed that evolution favored PD cooperation. This is because cooperation tends to increase reproductive potential. Thus, the strategies with the best chance of dominating were ethnocentrism and humanitarianism. However, ethnocentrism eventually surpasses humanitarianism for two possible reasons. First, when ethnocentric and humanitarian groups meet, ethnocentric agents tend to dominate along the group frontier by exploiting the humanitarian tendency to cooperate with everyone. Each of these two groups supports its own agents from behind the frontier lines, but at the frontier, ethnocentric agents develop somewhat better reproductive potentials than humanitarian agents do. Second, ethnocentrism may do better against selfishness than humanitarianism do, by virtue of not cooperating with selfish agents of different groups. Suppressing selfishness allows cooperation to flourish.

In evolutionary terms, the simulation demonstrated how genes look beyond their mortal carrier to their virtually immortal replicas in close relatives (Dawkins, 1989). Individual agents actually suffer by cooperating with their own kind, but the gene pool underlying the group is thereby strengthened relative to other groups.

All agents have to do to make this process work is to use a superficial marker (tag) to decide whether to cooperate with neighboring agents. Although these tags are initially meaningless and arbitrary, through evolution they come to have predictive significance for whether it is worthwhile to cooperate with another agent. It is interesting that rather minimal cognition is required of these agents – only discrimination of agent types and conditional action based on that discrimination.

Because Hammond and Axelrod (2006) focused on stable-state evolutionary outcomes (reached in the last 100 of 2000 cycles), and compared across parameter values, they did not study earlier stages of evolution. The possibility of such early stages within the evolutionary process for a single set of parameters is the focus of the present work.

Stages of Evolution

Method

We ran 50 different simulated worlds in the same fashion as Hammond and Axelrod, but analyzed and tested genotypic strategies at every one of 1000 cycles. We decided to stop at 1000 cycles because strategy dominance was already well established by that point. Each simulation was started with a unique random seed so that each simulated world was a unique realization of a stochastic process.

After each cycle in each run, we tabulated the number of agents with each of the four genotypic strategies and performed chi-square tests to ascertain whether any one of the four strategies had attained dominance over the others. For example, world0 after 18 cycles had 7 selfish agents, 5 traitors, 4 ethnocentrics, and 14 humanitarians. In that case, an overall chi-square test revealed nonsignificance, $\chi^2(3) = 8.83$, less than the critical value of 11.35 required at the $p < .01$ level. Expected values in each cell were computed by dividing the total number of agents by 4. After 70 cycles in world0, there were 21 selfish agents, 42 traitors, 39 ethnocentrics, and 82 humanitarians. At this point, the overall chi-square test did reach significance, $\chi^2(3) = 43$, $p < .01$, indicating that the distribution of strategy genotypes differed from chance. We followed this significant overall test with a single *df* chi-square test contrasting the two most frequent strategies. At this cycle, this test was also significant, $\chi^2(1) = 12.90$, $p < .01$, indicating that there were more humanitarians than traitors. The critical value for this latter test is 6.64 at the $p < .01$ level.

Thus, for a strategy to be considered dominant at a given cycle, both chi-squares had to be significant, one indicating that the frequency distribution differed from chance, and another that the most frequent strategy was significantly more frequent than the next most frequent strategy.

We used a rather strict level for significance testing, $p < .01$, rather than the more commonly used .05. We were not overly concerned with the large number of tests performed because, once a strategy started to dominate, its progress accelerated.

Results

Figure 1 shows strategy frequencies for each of the 1000 evolutionary cycles of world0. In this world, humanitarian strategies were statistically dominant in cycles 66-377. Ethnocentric strategies first became dominant at cycle 403, but because of continued stiff competition from humanitarian strategies, did not consistently maintain that statistical dominance until cycle 503. Twenty-six of the 50 worlds that we ran showed this pattern of early statistical dominance by humanitarianism.

The other common evolutionary pattern was characterized by early competition between humanitarian and ethnocentric strategies, without any consistent

dominance, that then gave way to ethnocentric dominance. An example is provided in Figure 2, where ethnocentric strategies became statistically dominant for good at cycle 178. Before that, there were no reliable differences between the two top strategies. Selfish and traitorous strategies never dominated at any cycle in these 50 worlds.

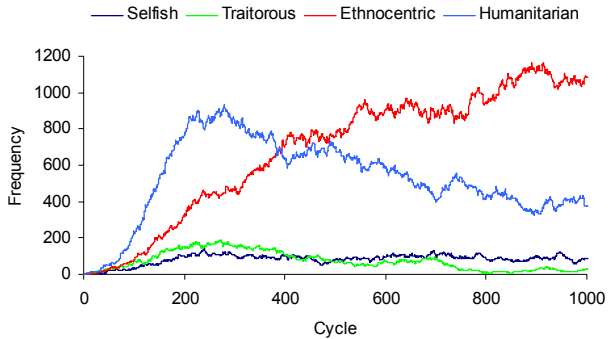


Figure 1: Strategy frequencies for world0. (Because it is difficult to distinguish four shades of gray, we present our figures in color. For those readers unable to access electronic copies of this paper or print colored copies, we could send a color hardcopy on request.)




Figure 2: Strategy frequencies for world6.

To get a more intuitive and more detailed understanding of the data, we also ran a few worlds yielding lattice plots of agent locations, strategies, and group tags every 20 cycles. Examples are presented in Figures 3 and 4 for world4, at 100 and 500 cycles, respectively. To better fit the formatting of this paper, we used a smaller, 25 x 25 lattice. The four tags are represented in these plots by different colors. Strategies are indicated by their first letter: s for selfish, t for traitorous, e for ethnocentric, and h for humanitarian.

At 100 cycles (Figure 3), there are 55 selfish agents, 15 traitorous agents, 44 ethnocentric agents, and 148 humanitarian agents. The corresponding proportions of these agent types are .21, .06, .17, and .56. The overall chi-square test of strategy frequencies was significant at 100 cycles, $\chi^2(3) = 152, p < .01$, as was that comparing the most frequent, humanitarian strategy to the next most frequent, selfish strategy, $\chi^2(1) = 43, p < .01$, confirming early humanitarian dominance. The emerging

correspondence between location, group, and strategy is already evident at 100 cycles. That is, adjacent agents come to share both strategy and group tag.




Figure 3: Strategies and groups in world4 at 100 cycles.

This correspondence is even stronger at 500 cycles (Figure 4), when this world is populated by 64 selfish agents, 7 traitors, 224 ethnocentrics, and 85 humanitarians, with corresponding proportions of .17, .02, .59, and .22. The two chi-square values are significant here as well, $\chi^2(3) = 268, p < .01$ and $\chi^2(1) = 63, p < .01$, but now reflect ethnocentric dominance.

Discussion

Cooperation can be healthy for evolution because it can increase reproductive potential, but only if many agents are doing it. Without generalized reciprocity, one-sided cooperation will increase others' fitness, but not that of one's own gene pool. For cooperation to flourish, selfishness needs to be suppressed in some way. Otherwise cooperative groups run the risk of being overrun by uncooperative free riders that leech off the generosity of other group members. In the case of our simulations, the effective mechanisms were in-group favoritism, keeping offspring close, and viscosity of the

environment. This allowed heavily cooperative strategies like ethnocentrism and humanitarianism to dominate.




Figure 4: Strategies and groups in world4 at 500 cycles.

These two strategies always competed early in evolution and in about half of the worlds there was an early stage in which humanitarian strategies were statistically dominant. Which of these two cooperative strategies dominates early may be largely a matter of chance, but ethnocentrism wins eventually, with about .75 of the agents. Important chance factors may include which agents immigrate, where they land, and how rapidly they reproduce, as well as mutation effects.

In contrast, selfish and traitorous strategies never dominated at any cycle in any of our worlds, although they might well do better under different constraints and parameter settings. Like the original Hammond and Axelrod (2006) simulation, this one shows that it is plausible that the pervasive human proclivity for ethnocentrism could have evolved.

Humanitarianism succeeds early, before groups have much contact, benefitting from cooperation and not being hurt much by interaction with defectors. But as the world fills in, humanitarian agents begin to lose PD games to defectors with consequent relative decrements in

reproductive potential, including especially losses to agents from continually expanding ethnocentric populations. Some of our current work is exploring different hypotheses about how and why humanitarians eventually lose out to ethnocentric populations.

We noticed in lattice plots that neighborhoods with considerable color diversity often included humanitarians. This is reasonable because humanitarian agents cooperate with, and thus enhance the reproductive potential of, adjacent agents regardless of their group and strategy. It suggests that humanitarianism may foster *multi-ethnic* regions. Our intuition is that this is the more likely direction of causal influence because agents with other strategies are likely to exploit the generosity of humanitarians thus diminishing the reproductive potential of humanitarian populations, if anything.

Why do both selfish and traitorous strategies fare so poorly in these simulations? The main factor here may be the inability of groups composed predominantly of either of these strategies to thrive. The idea is that groups of traitorous or selfish agents would be unlikely to gain any significant prominence because they do not cooperate with each other and thus have a lower chance of reproducing than humanitarian and ethnocentric groups. We can see from strategy frequency plots that even when the environment is still largely unoccupied and ripe for expansion (mostly before cycle 200), selfish and traitorous agents are fewer than ethnocentrics or humanitarians (see Figures 1 and 2 for examples). In current work, we are testing this and other hypotheses for the poor performance of selfish and traitorous strategies.

In all of our simulations, traitorous agents performed worse than selfish agents, even though traitors do cooperate, which we argue is beneficial to survival. Traitorous agents are handicapped not only by failing to cooperate with their own kind, but also by cooperating across group lines, particularly with ascendant groups of ethnocentrics. Whereas a selfish agent interacting across group lines with an ethnocentric agent will have an overall neutral outcome (no cooperation), a traitorous agent in the same situation will cooperate with the ethnocentric, and consequently fare worse in reproductive potential. This process can be seen in our data (Figures 1 and 2) where, as ethnocentrism becomes dominant, the number of traitorous agents begins to drop.

Although evolutionary change is commonly viewed as smooth and gradual, discontinuities have been noted, as for example in Gould and Eldredge's (1977) theory of punctuated equilibrium. Although their focus was on speciation through sexual reproduction, more abstractly there are some commonalities with our simulations in the notion of long periods of stasis punctuated by shorter bursts of rapid change, i.e., stages and transitions, respectively. There is some support for punctuated equilibrium in fossil records that fail to show the small, continuous changes predicted by gradual evolutionary theory (Somit & Peterson, 1992).

Such issues are rather difficult to study from historical fossil records alone because experimentation is impossible and there can be multiple stories fitting the often skimpy records that do exist. Agent-based simulations of evolution can overcome both of these hurdles – by running critical experiments that create their own complete evolutionary records. Like any other model or theory, such simulations typically abstract away a great many details present in the target phenomena. A good simulation will abstract away mostly irrelevant and unimportant details, leaving the essential, basic principles intact for systematic study. Still, it would be good to be able to test some of the predictions of evolutionary models in biological populations. This might be possible in species that evolve quickly in controlled, measurable laboratory settings.

The cognition employed by our agents may seem, and indeed is, fairly minimal. The key function in our computer code that implements a potential donor agent's determination to cooperate or defect is written as a nested *if* statement that takes account of whether the donor and recipient are members of the same group and what the donor's strategies are for in-group and out-group interactions. If the donor and recipient are in same group, then if the donor's strategy is to cooperate within its group, then the donor cooperates; otherwise it defects. But if the donor and recipient are in different groups, then if the donor cooperates outside of its group, then the donor cooperates; otherwise it defects. There is a separate function to compute whether the potential donor and recipient agents are from the same group, in terms of whether they possess the same tags.

Is such conditional action beyond the pale of what contemporary cognitive science regards as cognition? We think not, particularly in comparison to recent work in *minimal cognition*. In contrast to the dominant anthropocentric focus of most cognitive science, recent discoveries in microbiology indicate that cognition may be older and more fundamental than brains and nervous systems (van Duijn, Keijzer, & Franken, 2006). The fact that bacteria and other unicellular organisms are autonomous, social beings with the ability to differentiate aspects of the environment and respond accordingly suggests that cognition could have a primitive molecular basis that evolved into more complex neuronal cognition (DiPrimio, Müller, & Lengeler, 2000). The evolutionary study of such organisms might provide a more universal foundation for contemporary cognitive science. In any case, against this broad view of cognition, our agents can surely be viewed as cognitive.

A major omission from our simulations concerns learning, development, and various social influences. In future work, we plan to introduce such considerations into our models. Even so, agent-based studies of pure

evolution should continue to provide important bases of comparison in studying how learning and inheritance of learning strategies might interact with evolution.

In developmental psychology, nativist positions are often supported mainly by arguments that learning cannot explain the data (Fodor, 1980). There is rarely an analysis of, or direct evidence for, the evolution of genotypes. Simulations of evolution could compliment such nativist accounts by demonstrating the feasibility of evolutionary explanations and the effects of varying the basic parameters of evolution.

Nativist positions on development tend to minimize ontogenetic change while focusing on the constraints supplied by genotypic inheritance. Evolution, however, is also about change, albeit on a much slower time scale than ontogenetic developmental change, which is in turn on a slower scale than change brought about through learning. It is important to eventually understand change at all of these time scales and how they might affect each other.

Acknowledgments

This research is supported by a grant from the Natural Sciences and Engineering Research Council of Canada to the first author.

References

- Dawkins, R. (1989). *The selfish gene: New edition*. Oxford, UK: Oxford University Press.
- DiPrimio, F., Müller, B. S., & Lengeler, J. W. (2000). Minimal cognition in unicellular organisms. In *SAB2000 Proceedings Supplement: Sixth International Conference on Simulation of Adaptive Behavior: From Animals to Animats*. Honolulu: The International Society for Adaptive Behavior.
- Fodor, J. (1980). On the impossibility of learning "more powerful" structures. In M. Piatelli-Palmarini (Ed.), *The debate between Jean Piaget and Noam Chomsky* (pp. 142-152): Routledge & Kegan Paul.
- Gould, S. J., & Eldredge, N. (1977). Punctuated equilibria: The tempo and mode of evolution reconsidered. *Paleobiology*, 3, 115-151.
- Hammond, R. A., & Axelrod, R. (2006). The evolution of ethnocentrism. *Journal of Conflict Resolution*, 50, 926-936.
- Somit, A., & Peterson, S. A. (Eds.). (1992). *The dynamics of evolution*. Ithaca, NY: Cornell University Press.
- van Duijn, M., Keijzer, F., & Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*, 14, 157-170.