The Cognitive Cost of Ethnocentrism

Artem Kaznatcheev (artem.kaznatcheev@mail.mcgill.ca)

Department of Physics and School of Computer Science,
McGill University, 3600 University Street, Montreal, QC H3A 2T8 Canada

Abstract

Recent computational studies suggest that ethnocentrism, commonly thought to rely on complex social cognition, may arise through biological evolution in populations with minimal cognitive abilities. We use the methods of evolutionary game theory and computational modelling to examine the evolution of ethnocentrism. Since ethnocentric agents differentiate between in- and out-group partners, and adjust their behavior accordingly, they are more cognitively complex than humanitarian or selfish agents that always cooperate or defect, respectively. We associate a fitness cost with this complexity and test the robustness of ethnocentrism, concluding that ethnocentrism is not robust against increases in cost of cognition. Our model confirms that humanitarians are suppressed largely by ethnocentrics. Paradoxically, we observe that the proportion of cooperation is higher in worlds dominated by ethnocentrics. We conclude that suppressing free-riders, such as selfish and traitorous agents, allows ethnocentrics to maintain higher levels of cooperative interactions.

Keywords: Ethnocentrism; humanitarianism; cooperation; agent-based simulation; minimal cognition; Prisoner's Dilemma; evolution; free-rider-suppression hypothesis

Introduction

Seeing ones own group (in-group) as superior and other groups (out-groups) as inferior is a widespread syndrome of discriminatory behaviors (LeVine & Campbell, 1972). This perspective is associated with behavior including in-group favoritism (ethnocentrism) and out-group hostility (xenophobia) (Cashdan, 2001; Hewstone, Rubin, & Willis, 2002; Brown, 2004). Although the behavior is commonly thought to involve substantial cognitive ability (Sherif, 1966; LeVine & Campbell, 1972; Hewstone et al., 2002), extensive psychological evidence suggests that the presence of a strong in-group bias can be observed in individuals with minimal cognition and highly abstract social input (Tajfel, 1970; Tajfel, Billig, Bundy, & Flament, 1971). This is supported by observed ethnocentrism in human placenta (Haig, 1996), ants (Keller & Ross, 1998), and microbes (Lenski & Velicer, 2000), and suggests that ethnocentrism may have a basis in biological evolution. Cognitively, the ability to distinguish in- and out-group members and adjust behavior accordingly may be sufficient to foster this effect.

Recent computational studies (Hammond & Axelrod, 2006; Shultz, Hartshorn, & Hammond, 2008; Shultz, Hartshorn, & Kaznatcheev, 2009) have focused on the emergence of in-group favoritism through agent-based simulations of individuals with minimal cognitive ability. Agents interact via a one-time prisoners dilemma (PD) game that affects the reproductive potential of the participants. Agents can either defect against, or cooperate with, other in- or out-group agents, permitting four strategies: (1) a humanitarian strategy of universal cooperation, (2) an ethnocentric strategy of in- but not out-group cooperation, (3) a traitorous strategy of cooperation exclusively with the out-group, and (4) a selfish strategy of constant defection. Hammond and Axelrod (2006) showed that, after a transient period, ethnocentric agents dominate the population. Shultz et al. (2008) examined the transient period to uncover evidence for early competition between the ethnocentric and humanitarian strategies. More recently, Shultz et al. (2009) focused on explaining the mechanism behind ethnocentric dominance over humanitarians. In particular, they introduced the direct and free-rider-suppression hypotheses. The direct hypothesis is that ethnocentric clumps of agents directly suppress contacted clumps of humanitarian agents from different groups. The contrasting free-rider-suppression hypothesis is that ethnocentrics are more effective than humanitarians at suppressing groups of free riders — selfish and traitorous agents from the same group.

This paper extends beyond previous work by closely examining the cognitive mechanisms required for ethnocentrism. In particular, we measure the cost in fitness an agent is willing to pay in order to have the higher (yet still simple) cognitive processes required to discriminate between in- and out-groups. By varying the cost of cognition we also eliminate ethnocentric agents and confirm the direct hypothesis as the mechanism of ethnocentric dominance over humanitarians. Tracking the proportion of cooperation also reveals a novel and important role for the free-rider-suppression hypothesis — maintaining cooperative interactions.

Cognitive Complexity

An easy way to understand the complexity of reasoning carried out by simple abstract agents is to represent their decision procedure by finite state machines (FSMs). To use FSMs, it is important to understand what information agents receive and what actions they perform based on that information. During an interaction the agent receives some information about its partner and then makes a decision to cooperate or defect. In particular, the agent receives a signal S if the agent's partner is from the same abstract group (in-group) and N otherwise (out-group). In response, the agent outputs a D to defect in the PD interaction, and C to cooperate. Note that the agent does not receive any direct information

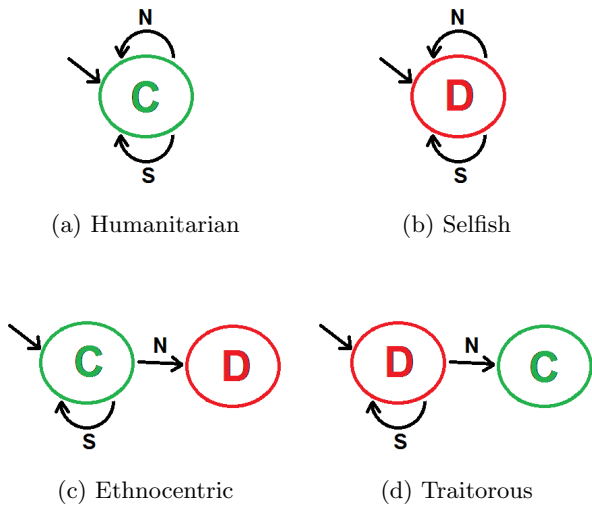(a) Humanitarian   (b) Selfish

(c) Ethnocentric   (d) Traitorous

Figure 1: Finite state machines representing the 4 possible strategies. Transitions are represented by arrows, with label S corresponding to an input of same-tag and N to different-tag. The agent's action is represented in the center of the state, C for cooperate and D for defect.

on its partner's strategy, and instead relies on potential correlations of a partner from the same group having the same strategy.

The four strategies are represented by FSMs in figure 1. Circles correspond to states, that are labeled by their output, either C or D. Arrows represent transitions, labeled by S or N. The initial state is signified by the smaller arrow with no pre-state. Humanitarian agents that always cooperate, and selfish agents that always defect are easy to represent. In particular, they are single state machines that output C or D (respectively) regardless of the input they receive. These two strategies are shown schematically in figures 1a and 1b. Ethnocentric and traitorous agents, however, require two states as shown in figures 1c and 1d. The extra state represents the greater complexity in making a decision based on input received, compared to not making a decision at all (the single state agents). Since implementing this rudimentary decision-making requires extra energy expenditure on the part of the agents, we represent this extra cost as a small fitness decrement k for ethnocentric and traitorous agents. This cost is especially important for studying the co-evolution of cooperation and ethnocentrism. To follow an ethnocentric strategy the agents have to invest a bit of their fitness into developing more sophisticated cognitive processes.

Method

Prisoner's Dilemma

In virtually any competitive social situation, interacting agents have a basic decision to make: cooperate

		Agent 2	
		C	D
Agent 1	C	$b - c$	$-c$
	D	b	0

Table 1: Payoff matrix for one agent (agent 1) interacting with another (agent 2).

with each other, or not. In evolutionary game theory such interactions are usually modelled by the Prisoner's Dilemma (PD). In PD, two agents independently decide whether to cooperate with or defect against the other. When cooperating, an agent pays a cost c to provide a greater benefit b to its partner. When defecting, an agent pays no cost and provides no benefit, but can still receive benefit if its partner cooperates. Table 1 shows the payoff for one agent (agent 1) interacting with another (agent 2). The payoff matrix reveals that an agent can always receive a higher payoff by defecting instead of cooperating. In game theoretic terminology, mutual defection is the Nash equilibrium. However, if both agents manage to coordinate cooperation, then they are both better of than if they had mutually defected — mutual cooperation Pareto dominates mutual defection. Due to this paradox, the PD game is regarded as a paradigmatic example of the problem of achieving mutual cooperation (Axelrod & Hamilton, 1981). The game provides a simple model of an environment where one action (defection) is better for the individual and the other (cooperation) is better for the population.

The simplest approach to studying the evolutionary dynamics of the PD is in a well mixed population. If agents are paired randomly from a mixed population and interaction results modify individual reproductive potential, eventually defectors will dominate the population. To allow cooperation to emerge it is essential to introduce positive correlations between the strategies of paired agents. To study the emergence of cooperation, researchers explore various ways to create these correlations. In our model we consider an interplay of spatial structure and arbitrary tags.

Model

Our model, and the Hammond and Axelrod (2006) model it is based on, expand beyond random interactions to facilitate the emergence of cooperation. Instead of randomly choosing interaction pairs, agents populate a toroidal square lattice (50 by 50 cells) and interact with their four adjacent neighbors. Each individual is simple, only perceiving whether it shares a common tag with neighbors (from a total of 4 tags), allowing for two interaction strategies: an in-group (igs) and an out-group (ogs) strategy. The four strategies are summarized in figure 1 and table 2. The outcomes of PD interactions

Name	igs	ogs	Figure
Humanitarian	C	C	1a
Ethnocentric	C	D	1c
Traitorous	D	C	1d
Selfish	D	D	1b

Table 2: The four possible strategies. The igs column correspond to the in-group strategy, ogs to out-group strategy.

(with $b = .025$ and $c = .01$) are added to the agents potential to reproduce (ptr, which is reset to .1 at the start of each cycle). At the end of a cycle, each agent has a chance equal to its ptr to clone itself (with a constant mutation rate (.005)) and a constant probability (.1) of dying. If an agent expires its location is vacated until habitation by a new agent. Regardless of the agent's survival, if the agent cloned itself the child is placed in one empty cell adjacent to the parent (potentially including the parent's cell if the parent expired after cloning itself). To start the world, and if the population ever reaches zero, the world is seeded with 80 individuals distributed randomly across the torus and uniformly across the 16 strains (4 possible tags, 2 possible igs, 2 possible ogs). The simulation runs for 3000 cycles.

To account for the potential cost of extra cognitive abilities we introduce a new parameter k. This parameter is varied in different simulations from 0 to .01 to show the impact of cognitive costs on the evolution of ethnocentrism. The effect of the cost of cognition is to lower the base ptr of ethnocentric and traitorous agents. In particular, at the start of every cycle the ptr is reset to .1 for humanitarian and selfish agents and to $.1 - k$ for ethnocentric and traitorous agents. The rest of the simulation is unmodified. By adjusting k we can quantify how much the extra cognitive powers of ethnocentrics are worth in terms of the currency of evolution — reproductive fitness.

During the simulation we collect two primary types of data. We record the distribution of agents by strategy, and track the number of cooperative and non-cooperative interactions. If an agent chooses to cooperate during its interaction, we increment the number of cooperations for the cycle. The proportion of cooperation is then the number of cooperations divided by twice the number of interactions (to account for both agents having to make a decision during each interaction).When comparing simulations with different values of k we take the mean data from the last 500 cycles, since the dynamics stabilize by then. To account for the stochastic nature of our model, we present all results with standard error from averaging over 30 worlds (a world is a single instance of the simulation with specific parameter setting and initial random seed).


Figure 2: Number of ethnocentric and humanitarian agents vs. cost of cognition. The points represent the mean number of agents over the last 500 cycles of the simulations: red is ethnocentric agents, blue is humanitarian. The error bars represent standard error from averaging over 30 different worlds.

Results

Our primary result is the variation in the number of ethnocentric and humanitarian agents and the proportion of cooperation for different values of the cost of cognition, k. As in previous studies (Hammond & Axelrod, 2006; Shultz et al., 2008, 2009) the number of selfish and traitorous agents is negligible, and hence we concentrate our presentation on the number of humanitarians and ethnocentrics in the population. Figure 2 shows the number of ethnocentrics (in red) and humanitarians (in blue) in simulations with increasing costs of cognition. Unsurprisingly, as the cost of cognition increases the number of ethnocentric agents decreases, and humanitarians take their place. This is consistent with the assessment of Shultz et al. (2009) that humanitarians are directly suppressed by ethnocentrics. A surprising result, is how low the cost of cognition ($.004 \leq k \leq .0045$) is at the point where the humanitarian strategy becomes fitter. In particular, this transition point is more than an order of magnitude lower than the default ptr (.1) and less than half of the cost of cooperation (.01). The third interesting result, is the quick phase transition from ethnocentric to humanitarian dominance. With $k \leq .0035$ the number of humanitarian agents is relatively stable around 200, from $k = .004$ to $k = .0065$ we observe quick change, and then for $k \geq .007$ the humanitarian population stabilizes around 1000 individuals. Together, these results suggest that ethnocentrism is not very robust against variance in the cost of cognition. In particular, for widespread ethnocentrism to emerge, the cost of differentiating between in- and out-groups needs to be

Figure 3: Proportion of cooperative interactions vs. cost of cognition. The points represent the mean proportion of cooperative interactions over to last 500 cycles. The error bars represent standard error from averaging over 30 different worlds.

extremely low in comparison to other relevant parameters (default ptr, cost and benefit of cooperation, etc.).

In figure 3 we show the proportion of cooperative interactions as the cost of cognition increases. Although the change in cooperation is not as drastic as the changes in strategy distribution, it is still statistically significant. In particular, we observe the same phase transition, with proportion of cooperative interactions stable around .955 while $k \leq .0035$ and stable around .915 when $k \geq .0065$. The counter-intuitive result in figure 3 is that as humanitarian agents start to dominate the population, the proportion of cooperation decreases. This raises the important question of what is more important: overall cooperation or the individual fairness of predominantly humanitarian agents?

For contrast, we also present two figures of strategy distribution by cycle. In figure 4a we show a cost of cognition $k = .002$, a bit before the phase transition. In figure 4b we examine a point after the phase transition, $k = .007$. The drastic change from ethnocentric (green) dominance to humanitarian (blue) dominance is self-evident. Further, in the humanitarian-dominated world of $k = .007$ selfish (red) agents perform around 5 times better than in the ethnocentric dominated world. This supports the intuition that ethnocentrics are better at suppressing selfish agents.

For other nearby choices of basic parameters (b, c, base ptr, and death rate) the qualitative results are similar, although the exact quantitative aspects change. As in reports (Shultz et al., 2008, 2009), we omit them for brevity.

(a) $k = .002$

(b) $k = .007$

Figure 4: Number of agents of various strategies vs. evolutionary cycle. The lines represent the number of agents of each strategy: blue — humanitarian; green — ethnocentric; yellow — traitorous; red — selfish. The width of the line corresponds to standard error from averaging 30 different worlds. The two figures correspond to different costs of cognition, k.

Discussion

The relatively low cost of cognition (around $.004 \leq k \leq .0045$) required to transition from ethnocentric to humanitarian dominance suggests that ethnocentrism is not very robust. In particular, the emergence of ethnocentric cooperation is unlikely to have caused significant investment of fitness in cognitive development. Alternatively, the mechanisms for differentiating between in- and out-groups and making basic decisions would need to be already in place by other means, and are unlikely to have co-evolved with cooperation. Making the distinction between in- and out-groups does not require fitness investment for humans, but for more rudimentary organism, it is likely that it would. Although the cognitive abilities required for ethnocentrism are as simple as distinguishing in- and out-groups, this simplicity can be deceiving. Our results stress that for ethnocentrism to evolve, these simple cognitive abilities must also be extremely cheap in terms of fitness invested.

The results in figure 3 suggest that displacing ethnocentric agents by humanitarian ones can lead to a decrease in overall proportion of cooperative interactions. In particular, it is important to reexamine the negative perception of ethnocentrism. Although unfair from the individual point of view, ethnocentrism might be essential to sustain the levels of cooperation required for complex structures such as multi-cellular organisms or human society. Thus, a tempting answer to the question of Shultz et al. (2009): "why is ethnocentrism more common than humanitarianism?" is that humanitarianism cannot maintain as high levels of cooperation.

A further connection to previous work (Shultz et al., 2009) is a reevaluation of the direct and free-rider-suppression hypothesis. Although Shultz et al. (2009) ruled out the free-rider-suppression hypothesis in favor of the direct hypothesis, they did not examine the proportion of cooperation. The direct hypothesis provides a good explanation of why ethnocentrics dominate humanitarians, but the free-rider-suppression hypothesis explains the increased levels of cooperation in largely ethnocentric populations. When humanitarians replace ethnocentric as the dominant strategy, significantly higher levels of selfish agents evolve in the population. The decrease in cooperation caused by higher levels of selfish agents exceeds the increase in cooperation caused by humanitarians cooperating across groups. This results in an overall reduction in the cooperative interactions. Thus, ethnocentric agents ability to better suppress free-riders is important for maintaining higher levels of cooperative behavior.

Although the decision making employed by our abstract agents is extremely simple, it is not beyond the scope of what contemporary cognitive science regards as cognition. Our research explores rudimentary cognition in a social and evolutionary context. In particular we hope that this paper highlights the importance of considering possible fitness investment in even the simplest forms of cognition. By exploring further we hope to gain a better understanding of the evolution and potential social effects of simple information processing.

References

Axelrod, R., & Hamilton, W. (1981). The evolution of cooperation. *Science*, *211*, 1390-1396.

Brown, D. E. (2004). Human universals, human nature, and human culture. *Daedalus*, *133 (4)*, 47-54.

Cashdan, E. (2001). Ethnocentrism and xenophobia: A cross-cultural study. *Current Anthropology*, *42*, 760-765.

Haig, D. (1996). Gestational drive and the green-bearded placenta. *Proc. Natl. Acad. Sci. USA*, *93*, 6547-6551.

Hammond, R., & Axelrod, R. (2006). The evolution of ethnocentrism. *Journal of Conflict Resolution*, *50*, 926-936.

Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology*, *53*, 575-604.

Keller, L., & Ross, K. (1998). Selfish genes: a green beard in the red fire ant. *Nature*, *394*, 573-575.

Lenski, R., & Velicer, G. (2000). Games microbes play. *Selection*, *1*, 51-57.

LeVine, R. A., & Campbell, D. T. (1972). Ethnocentrism. *New York: John Wiley*.

Sherif, M. (1966). Group conflict and co-operation: Their social psychology. *London: Routledge Kegan Paul*.

Shultz, T. R., Hartshorn, M., & Hammond, R. A. (2008). Stages in the evolution of ethnocentrism. In B. Love, K. McRae, & S. V.M. (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (p. 1244-1249). Austin, TX: Cognitive Science Society.

Shultz, T. R., Hartshorn, M., & Kaznatcheev, A. (2009). Why is ethnocentrism more common than humanitarianism? In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the congnitive science society* (p. 2100-2105). Austin, TX: Cognitive Science Society.

Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific America*, *223*, 96-102.

Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology*, *1*, 149-178.