The Dangers of Ethnocentrism

Giangiacomo Bravo
Department of Social Studies, Linnaeus University
Email: giangiacomo.bravo@lnu.se

ABSTRACT

Humans often alter their behavior depending on the opponent's group membership, with positive (e.g., support of same-group members) or negative (e.g., stereotyping, oppression, genocide) consequences. An influential model developed by Hammond and Axelrod highlighted the emergence of macro-level "ethnocentric cooperation" from the aggregation of micro-level interactions based on arbitrary tags signaling group membership. In this paper, we replicated this model and extended it to allow a wider array of possible agents' behaviors, including the possibility of harming others. This allowed us to check whether and under which conditions xenophobia can emerge beside or in alternative to ethnocentrism.

INTRODUCTION

Ethnocentrism, more properly known as intergroup bias, is the widespread attitude of humans to change their behavior depending on the opponent's group membership. It can have positive implications, such as helping same-group members, but also lead to negative behaviors towards out-group members—ranging from prejudice and stereotyping, to oppression, and genocide—which are often popularized as examples of xenophobia (Hewstone et al. 2002). Group membership is usually built upon symbolic markers—easily observable characteristics such as language, clothing, skin color, etc.—that may dramatically influence people's behavior (Kurzban et al. 2001; Richerson and Boyd 2001). When group boundaries are dependent on geographical proximity, these symbolic aspects of social identity may also determine spatial segregation between ethnic groups (e.g., Musterd 2005).

Some years ago, Hammond and Axelrod (2006b) developed an influential agent-based model (hereafter HA model) that was aimed at cast new light on these issues. They suggested to look at the emergence of macro-level "ethnocentric cooperation" from the aggregation of micro-level interactions based on a standard Prisoner's dilemma game (hereafter PD game). Their definition of ethnocentric cooperation was centered on in-group favoritism based on the interplay between arbitrary "tags", signaling group membership, and local interaction. The research showed that, once combined with "population viscosity", tags were sufficient to determine "high levels of individually costly cooperation with only minimal cognitive requirements and in the absence of other, more complex mechanisms", such as the action of a central authority or reciprocity-based strategies (Hammond and Axelrod 2006b, 932).

More recently, Hartshorn et al. (2013) replicated this model and suggested to reconsider Hammond and Axelrod's optimism over the positive effect of ethnocentrism in favoring cooperation. They argued that the most likely alternative of Hammond and Axelrod's ethnocentric cooperation is not the lack of cooperation but "humanitarianism", i.e., the willingness to help others independently on their tags. Although important, this study had several shortcomings. In their analysis of the mechanisms through which ethnocentrism develops, Hartshorn and colleagues failed to highlight that the conditions leading to ethnocentric cooperation were the same producing universal cooperation when the possibility of identifying group membership was ruled out (see Hartshorn et al. 2013, Tab. 3). In other words, the high cooperation levels in the HA model could be more the result of local agent interaction than a by-product of tags (see also Jansson 2013).

A second crucial issue is that, while Hammond and Axelrod's explored only the positive side of intergroup-bias, the reality is full of examples of its dark side: xenophobia. Although xenophobia is not an inevitable result of ethnocentrism (e.g., Cashdan 2001), the two features often go side by side (Brewer 2001; Hewstone et al. 2002). Therefore, the fact that the same model leading to ethnocentrism could also explain the emergence of xenophobic behaviors, if only the authors would have allowed a wider range of strategies to evolve, requires a further in-depth analysis, which is presented in this paper. First, a further replication of the HA model will be developed to disentangle the conditions leading to the evolution of ethnocentric behaviors from the ones favoring to unconditional altruism (Study 1). Then the assumption limiting the possible actions to positive behaviors will be relaxed and agents will be allowed to actively harm others (Study 2). This will allow to check whether and under which conditions xenophobia can emerge beside or in alternative to ethnocentrism.

STUDY 1

In this study, we replicated the HA model and extended the analysis originally done by the authors to include no-tag and random-location conditions. Note that these possibility were studied by Hammond and Axelrod using a slightly different model in a paper published in the *Theoretical Population Biology* (hereafter TPB) journal (Hammond and Axelrod 2006a), but neither presented nor discussed in their *Journal of Conflict Resolution* (hereafter JCR) article (Hammond and Axelrod 2006b).

In the HA model, agents are located on a space formed by a $N \times N$ toroidal lattice, meaning that the borders wrap around such that each site on the lattice has exactly four neighbors, i.e., the agents located on the four adjacent North, South, East and West locations. Agents can only interact with their neighbors. In the interaction phase, they decide whether to pay a cost (of giving help) $c > 0$ to give a benefit (of receiving help) $b > c$ to their neighbors, taking a

Proceedings 28th European Conference on Modelling and Simulation ©ECMS Flaminio Squazzoni, Fabio Baronio, Claudia Archetti, Marco Castellani (Editors)
ISBN: 978-0-9564944-8-1 / ISBN: 978-0-9564944-9-8 (CD)

separate decision for each neighbor. Note that this structure of the game is equivalent to a standard PD game. Therefore, we have labeled the two possible actions of giving or not giving help as cooperation (C) and defection (D) respectively. Costs ad benefits are then subtracted/added to the basic "potential to reproduce" (PTR) of each agent, i.e., its probability to produce an offspring in a given period, which is set at a fixed initial at the beginning of the period and is then increased by b each time the agent receives help while it is decreased by c each time it helps.

Each agent i holds an arbitrary tag that was established at its birth and is kept fixed throughout its life. The agent's strategy was determined by the couple (S_i^s, S_i^d), where $S_i^s, S_i^d \in \{C, D\}$ with C meaning to cooperate (or "help") the neighbor and D to defect (or "not help") it. Agents with $S_i^s = C$ will cooperate with neighbors holding the same tag, while agents with $S_i^d = C$ will cooperate with neighbors holding a different tag; $S_i^s, S_i^d = D$ correspond to defection with agents holding the same and a different tag respectively. Agents can hence follow one these different strategies: the "ethnocentric" (C,D), the "altruistic" (C,C), the "selfish" (D,D) and the "traitor" (D,C) one.

The simulation runs for T periods. At the beginning of each run, the space is empty. Then in each period the following four stages occur.

1. An "immigrant"—i.e., a new agent—is created in an empty cell with random strategy and tag.
2. Agents' PTRs are reset to their basic values, then all interactions occur.
3. Agents reproduce with a probability equal to their PTR. Newborns are placed in one of the four neighboring sites, which implies that agents can only reproduce if they have less than four neighbors. Newborns inherit their parents' characters (S_i^s, S_i^d and tag). However, these can randomly change with probability m.
4. Each agent dies with probability d.

As Hammond and Axelrod (2006b) found that their model results were rather robust to parameter changes, we will present here only results following their standard parameter definition, namely $N = 50$, $T = 2000$, basic PTR $= 0.12$, $c = 0.01$, $b = 0.03$, $m = 0.005$, $d = 0.1$. The model was implemented in NetLogo (Wilensky 1999).

In order to discriminate between the effect of tag-based strategies and the spatial distribution of agents in the HA model, we proceeded in two steps. First, we compared the outcome of the original configuration with the one obtained by a model where we excluded the tag effect, i.e., where all agents hold the same tag and $S_i^s = S_i^d$. This means that only two strategies are possible: the standard selfish (C,C) and altruistic (D,D) ones. Second, we relaxed the local reproduction rule, i.e., instead of being placed in one of the neighboring sites, newborns can be assigned to any empty patch of the simulated space. This means that four conditions are possible: tags - adjacent location (T-A); no tags - adjacent location (N-A); tags - random location (T-R); no tags - random location (N-R). Being based on a 2×2 full factorial design, this analysis allowed to estimate not only the main effects of tags and localization (and to compare their relative strength) but also the effect of the interaction between the two variables.

Results

For all parameter configurations, we performed 50 runs of the model recording the distribution of agent strategies and their moves in each period. The proportion of cooperative moves in the last 100 periods of the T-A condition, when the system settled to its equilibrium, was 0.80 ± 0.0003, with the ethnocentric strategy followed by 84% (Tab. I). This was even higher than Hammond and Axelrod (2006b) findings (76% of ethnocentric agents), but it should be noted that their sample was limited to 10 repetition, so that this difference could be due merely by chance. On the other hand, this figure is somewhat lower than the 89% one presented in the TPB paper (Hammond and Axelrod 2006a, Tab. 1). Similarly, the N-A condition led to a high proportion of cooperative moves (0.86 ± 0.0003), with a large dominance of altruistic strategies (86%).

The outcome dramatically changed when the adjacent location rule was relaxed. In the T-R condition the proportion of cooperative moves was only 0.03 ± 0.0003, while it was 0.02 ± 0.0001 in the N-R condition. In both cases, selfish strategies largely proliferated. Note also that, given that the number of offspring was no longer limited by the available space in the neighborhood of existing agents, the average number of agents was higher in both random location conditions (Tab. I). In all cases, the dynamics of the simulation led to an early establishment of one of the strategies with limited changes afterwards (Fig. 1).

Regressing the simulation conditions on the proportion of cooperative moves led to even clearer results. OLS estimates of a model considering both the main and interaction effects of the two variables showed that, over a background of full defection (except for the randomness introduced by the mutation rate), the introduction of tags increased the proportion of cooperative moves by 0.015 ± 0.003 ($t = 7.88$, $p < 0.001$) while introducing the adjacent location rule increased it by more than 50 times, namely by 0.847 ± 0.003 ($t = 283.70$, $p < 0.001$). Note also that the interaction term estimate had a negative sign (-0.084 ± 0.004, $t = -19.78$, $p < 0.001$), meaning that tags actually *reduced* the positive effect of adjacent location on cooperation.

Discussion of Study 1 results

Our results showed that tag-based strategies do not favor cooperation *per se*. On the contrary, by discriminating between agents holding different tags, they reduce the potential level of cooperation ideally favored by the adjacent location condition. This is clear contrasting the B and D panels in Figure 2. In the B one, whit no tags allowed, altruists form large cluster with only small marginalized groups of selfish agents. On the contrary, in the D panel almost all borders between clusters of agents holding different tags are marked by non-cooperative moves, mainly performed by ethnocentric agents, leading to overall higher defection levels. Note also that the introduction of tags in the random-location condition did not significantly improved cooperation. This because in well mixed populations the probability of meeting an agent with a different tag is higher than the one of meeting an agent with the same tag.

Note that, as highlighted by Nowak et al. (1994), spatial proximity allows*per se* selfish and altruistic strategies to

	Adjacent location				Random location			
	Tags		No tags		Tags		No tags	
	mean	sd	mean	sd	mean	sd	mean	sd
(C,C) *altruist*	0.096	0.032	0.857	0.028	0.005	0.002	0.016	0.005
(C,D) *ethnocentric*	0.837	0.037	0.000	0.000	0.056	0.023	0.000	0.000
(D,C) *traitor*	0.016	0.008	0.000	0.000	0.017	0.005	0.000	0.000
(D,D) *selfish*	0.052	0.016	0.143	0.028	0.922	0.025	0.984	0.005
Average S_i^s	0.932	0.018	0.857	0.028	0.061	0.024	0.016	0.005
Average S_d^s	0.111	0.035	0.857	0.028	0.022	0.006	0.016	0.005
Cooperative moves	0.795	0.021	0.864	0.027	0.032	0.008	0.017	0.005
Number of agents	1583.132	44.250	1609.550	46.439	2272.999	0.042	2273.000	0.000

TABLE I: Study 1 results. The table reports statistics for the last 100 periods of the simulations.

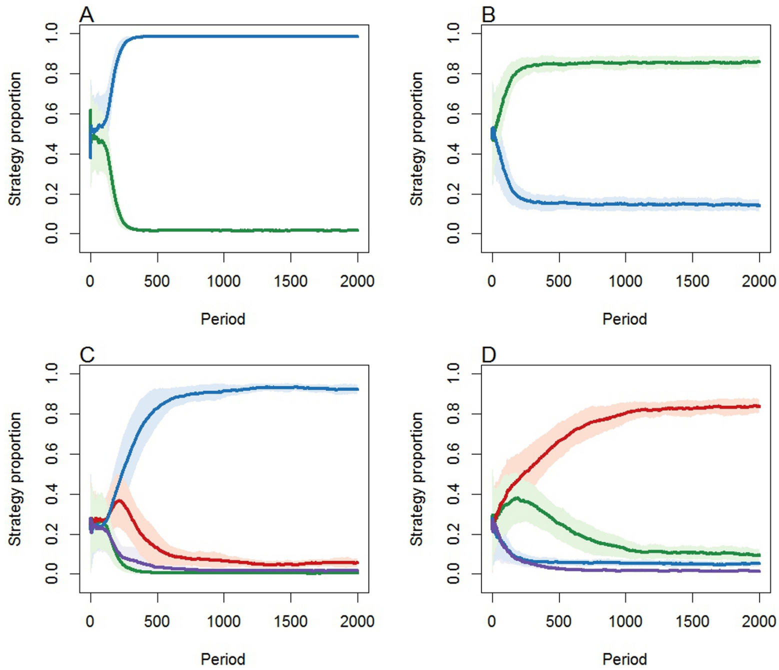

Fig. 1. Dynamics of the HA model. The figure presents the average proportion of each strategy for all model runs. Panels: (A) no tags - random location; (B) no tags - adjacent location; (C) tags - random location; (D) tags - adjacent location. Color legend: green = altruist; red = ethnocentric; blue = selfish; purple = traitor. Colored bands indicate standard deviations.

coexist, with cooperation spreading under a relatively large set of conditions (depending on the b/c ratio and on the probability that a newborn is placed in its parent's neighborhood). Since these conditions are met in the N-A case, the large share of altruist agents at the end of the simulation are fully consistent with this (Fig. 1B and 2B). More generally, as highlighted also by Axelrod et al. (2004), the evolution of altruism under these conditions can be explained mainly trough "inclusive fitness" (Hamilton 1964), which states that natural selection will favor the spread of cooperative behaviors as long as $rb-c>0$, with r representing the coefficient of relatedness between two agents (obviously one between a parent and its offspring in our case, except for mutation). Given the current parameter configuration, this means that, in absence of tags, altruism will spread if altruists have an average number of cooperative neighbors equal or greater to two: a condition that is respected under the adjacent location, but not under the random location condition.

Hammond and Axelrod (2006a) actually acknowledged this in their TPB paper, although arguing that the joint mechanism of tags and population viscosity "vastly increases the range of environments in which contingent altruism can evolve". Moreover, in their JCR paper they do not even mention the possibility that cooperation evolves as an effect of the spatial proximity in absence of tags (Hammond and Axelrod 2006b), highlighting the role of ethnocentrism in maintaining high cooperation levels in structured populations. However, the closest competitor for ethnocentrism in viscous populations is not selfishness but altruism, with the former being the main cause of the decline of the latter (Hartshorn et al. 2013). This means that, ethnocentrism should be seen more as a factor *reducing* cooperation rather than something contributing to its development. To sum up, Hammond and Axelrod (2006b) paper had the clear merit to show the mechanisms through which ethnocentrism is likely to emerge and persist in structured populations. At the same time, their claim that " in-group favoritism can be an undemanding yet powerful mechanism for supporting high levels of individually costly cooperation" should be reconsidered as the same conditions they studied are likely to produce cooperation also in absence of ethnocentric strategies.

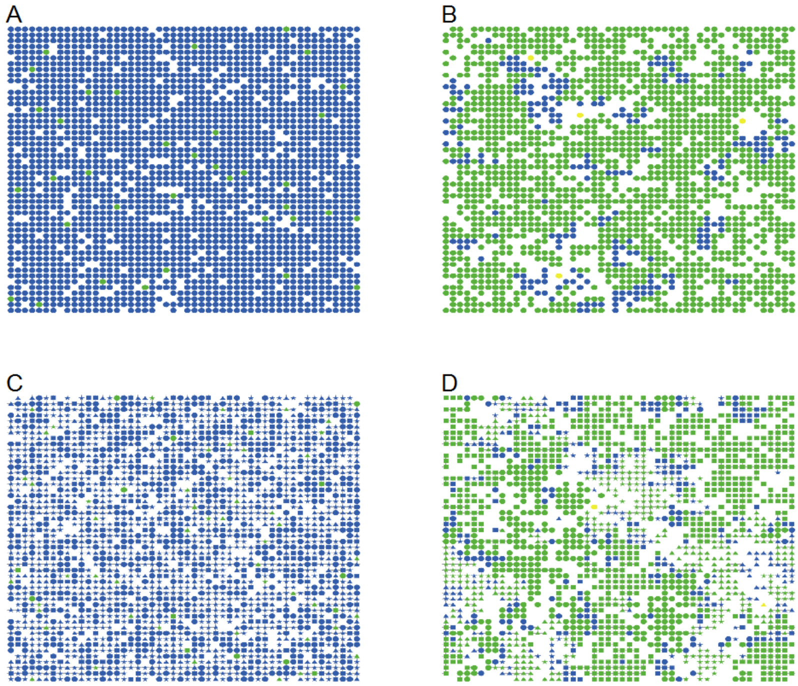

Fig. 2. View of the interaction space at the end of a typical run before reproduction and death under all parameter configurations. Note that only the last move of each agent is shown. Panels: (A) no tags - random location; (B) no tags - adjacent location; (C) tags - random location; (D) tags - adjacent location. Agents' tags are represented by symbols, agents moves by colors: green = C and blue = D. Yellow agents are isolated ones that did not interact in the current period.

Study 2

The effect of intergroup bias is not always limited to restraining to cooperate with out-group individuals, but can also take the form of actively harming who is not recognized as a group member (Brewer 2001; Hewstone et al. 2002). This is especially relevant for the evolution of human societies, as warfare among tribes was widespread among hunter gatherers, and ethnic clashes are widespread also today(e.g., Bowles 2012; Keeley 1996; Richerson and Boyd 2001).

A straightforward extension of the HA model is to allow agents not only to benefit, but also to impose a cost on their neighbors. This represents any action harming other people, independently of its specific content (physical attack, limits to personal freedom, etc.). Formally, let us think of a game where, besides choosing as above between helping (C) or not helping (D), players have the possibility to bear a cost c to *reduce* by b their opponents' PTR, with $b > c > 0$ as before. We called this third option "harming" (H), as it implies an actual reduction the reproduction probability of the opponent. The payoff table of the harming game is presented in Figure 3. Note that the PD game used in Study 1 is a sub-game (given by the four upper-left cells) of the harming one. A the same time, the only Nash equilibrium of the new game is mutual defection, just as in the PD one.

In the harming game model (hereafter HG model), a player's strategy is defined, as before, by the couple (S_i^s, S_i^d). However, now S_i^s and S_i^d take the values $\{C, D, H\}$ and nine possible strategies exists. Among them, especially important is the (C,H) one, which makes agents both to help those sharing the same tag and to harm agents with different tags. We called this new strategy "xenophobic" (see Hewstone et al. 2002), while (C,D), (C,C), (D,D) and (D,C) are called ethnocentric, altruistic, selfish and traitor respectively, as in the previous study. Except this, everything was equal to Study 1.

Results

We performed 50 runs of the HG model for each of parameter configuration (the same used in the first study), recording in every period the distribution of agent strategies and their moves. As before, the proportion of cooperative moves was considerably higher in the adjacent than in the random location conditions, with the introduction of tags making little difference (Tab. II). However, the distribution of agent strategies was different in the two conditions. While unconditional cooperation dominated in the N-A condition, the modal strategy was the xenophobic one in the T-A condition. Note also that the (D,H) strategy was the only challenging the dominance of altruism in the T-R condition (Tab. II and Fig. 4).

As before, OLS estimates showed that, starting from a situation of universal defection (except for the randomness introduced by mutation), the introduction of tags increased the cooperation proportion only by 0.006 ± 0.0004 ($t = 14.79$, $p < 0.001$) while the introduction of adjacent location increased it by 0.841 ± 0.0004 ($t = 2041.07$, $p < 0.001$). Moreover, the interaction of tags with adjacent location produced a small but significant decrease of cooperation (-0.006 ± 0.0005, $t = -10.14$, $p < 0.001$).

The proportion of harming moves was relatively low under all conditions even if it significantly increased, up to over 5% of all moves, when both tags and adjacent location were present (II). OLS estimates showed that the proportion of attacks increased by 0.009 ± 0.0002 when tags were

		Player 2		
		Cooperate	Defect	Harm
Player 1	Cooperate	$b-c, b-c$	$-c, b$	$-b-c, b-c$
	Defect	$b, -c$	$0, 0$	$-b, -c$
	Harm	$b-c, -b-c$	$-c, -b$	$-b-c, -b-c$

Fig. 3. The harming game

	Adjacent location				Random location			
	Tags		No tags		Tags		No tags	
	mean	sd	mean	sd	mean	sd	mean	sd
(C,C) *altruist*	0.089	0.039	0.845	0.032	0.002	0.002	0.011	0.003
(C,D) *ethnocentric*	0.340	0.132	0.000	0.000	0.022	0.019	0.000	0.000
(C,H) *xenophobic*	0.500	0.129	0.000	0.000	0.003	0.002	0.000	0.000
(D,C) *traitor*	0.009	0.007	0.000	0.000	0.057	0.079	0.000	0.000
(D,D) *selfish*	0.021	0.011	0.137	0.031	0.740	0.282	0.979	0.005
(D,H)	0.027	0.013	0.000	0.000	0.151	0.236	0.000	0.000
(H,C)	0.002	0.002	0.000	0.000	0.002	0.002	0.000	0.000
(H,D)	0.005	0.003	0.000	0.000	0.021	0.015	0.000	0.000
(H,H)	0.006	0.003	0.017	0.005	0.003	0.002	0.010	0.004
Average S_i^s	0.917	0.019	0.828	0.033	0.002	0.017	0.001	0.005
Average S_d^s	–0.432	0.139	0.828	0.033	–0.096	0.186	0.001	0.005
Cooperative moves	0.852	0.026	0.852	0.031	0.017	0.006	0.011	0.003
Harming moves	0.054	0.014	0.016	0.005	0.019	0.005	0.010	0.004
N. of agents	1733.168	32.072	1766.670	29.907	2272.990	0.325	2272.996	0.135

TABLE II: Summary of study 2 results. The table reports statistics for the last 100 periods of the simulations.

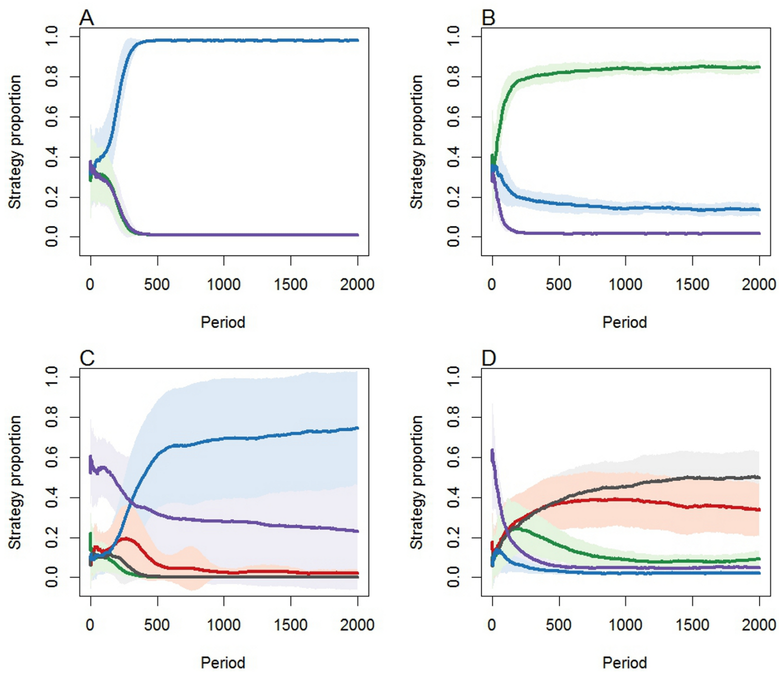

Fig. 4. Dynamics of the harming game. The figure shows the average proportion of each strategy for all model runs. Panels: (A) no tags - random location; (B) no tags - adjacent location; (C) tags - random location; (D) tags - adjacent location. Color legend: green = altruist; red = ethnocentric; gray = xenophobic; blue = selfish; purple = all other strategies. Colored bands indicate standard deviations.

used ($t = 51.67$, $p < 0.001$), by 0.006 ± 0.0002 when adjacent location was introduced ($t = 37.58$, $p = 0.001$), and by 0.029 ± 0.0002 when both factor were at work ($t = 124.73$, $p < 0.001$).

Discussion of Study 2 results

In the HG model, adjacent location was the factor affecting cooperation the most, just as in the first study. Harming choices were infrequent. Nevertheless, important differences exist between the four conditions. Notably, the fact that harming choices occurred only slightly more often in the in the T-A than in the other conditions hinders some crucial differences. In the T-A condition, harming occurred infrequently only because it was limited by the presence of well defined borders between the tag-groups (Fig. 5D). On the other hand, the overall harming potential (measured as the proportion of agents with $S_i^s = H$ and, espe-

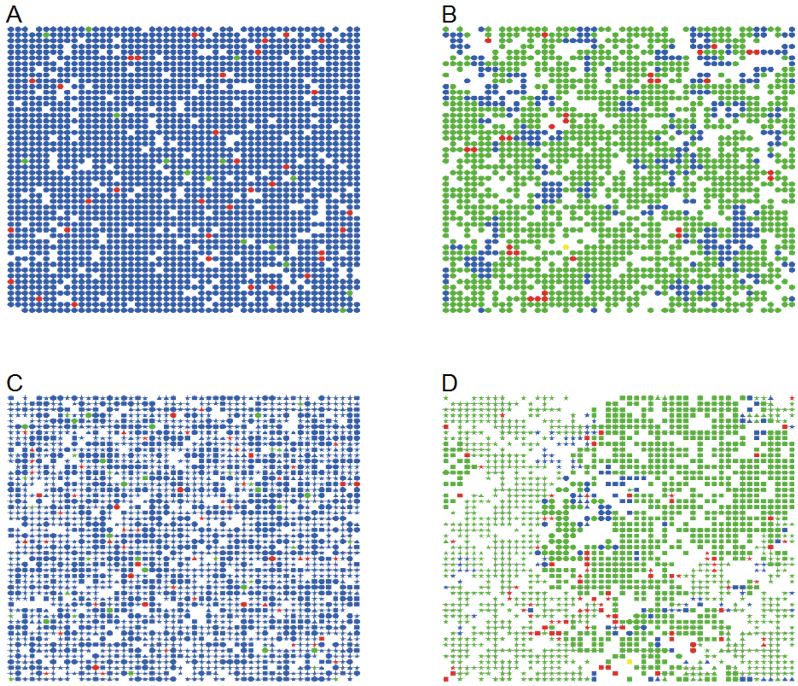

Fig. 5. View of the interaction space at the end of a typical run before reproduction and death under all parameter configurations. Note that only the last move of each agent is presented. Panels: (A) no tags - random location; (B) no tags - adjacent location; (C) tags - random location; (D) tags - adjacent location. Agents' tags are represented by symbols, agents moves by colors: green = C, blue = D, and red = H. Yellow agents are isolated ones that did not interact in the current period.

cially, $S_i^d = H$) was much higher in this condition than in any other one. In other words, while the harming potential of the no-tag conditions is fully exploited, any factor increasing the mixing of agents in the tag ones would have fostered a boost of harming behaviors. This is especially true in the T-A condition, where large cluster of same-tag agents existed that were composed by a majority of agents willing to harm out-group members.

The effects of the harming potential can be revealed by artificially modifying the system once it reached the equilibrium. For instance, changing the agent tags in the T-A condition such that large similar-tag groups were no longer present systematically produced a sudden burst of harming choices, up to 50% of total moves. A similar effect, was also present in the T-R condition, although in the long run this intervention tended to favor the evolution of the selfish strategy at the expenses of the harming ones, (C,H) and (D,H).

Harming strategies evolved mainly because they created extra room for the reproduction of kins (same-tag, same-strategy agents), which means that the spreading of the harming potential depended on the same adjacent location assumption that favors altruism in absence of tags. The effect of this spatial assumption (which may not necessarily hold in the real world) was especially relevant in the T-A condition, where the presence of large same-tag groups allowed borderline agents to afford the extra cost due to their harming actions towards out-group members by means of the cooperation of same-tag agents. In other words, given a sufficiently high segregation, the cooperation of the ethnocentric agents supports xenophobic behaviors, in a situation of constant between-group conflict which probably mimics human human evolution (Bowles 2012; Keeley 1996).

Finally, it is worth noting that, while in Study 1 ethnocentric cooperation was at least better than universal defection, this was no longer true when considering xenophobic strategies. Consider a simpler situation with only four agents, two of them holding tag A and two tag B. If everyone played selfishly they would obviously get zero in each period. If everyone used the xenophobic strategy, each of them would pay $3c$, once for helping the same-time agent and twice for harming the different-tag agents, while it would receive b because of the help from the same-tag agent and $-2b$ because of the harming from the two different-tag agents. The resulting payoff is $-b-3c$, which is always smaller than zero. Therefore, the introduction of tags led to a situation that was worst not only than the one where cooperation dominates, but also than the universal defection one.

Conclusions

Despite the optimistic conclusions of Hammond and Axelrod (2006b), we have shown that the development of ethnocentric strategies might not only limit the potential of cooperation in structured populations, but should be considered as a serious danger due to its "natural" coexistence of this behavior with less peaceful strategies. When the ban of actively harm out-group members is removed, a significant room for the development of xenophobia opens. Moreover, ethnocentric and xenophobic agents easily coexist and mutually reinforce at the expenses of other groups.

The emergence of symbolic markers linked with group distinction should hence be seen more as a potential treat for the pacific coexistence among individual than as a strategy to increase cooperation in the system. Also spatial segregation looks like a double-edged sword. On the one hand, it can lead to high levels of universal cooperation in absence of tags. On the other, it easily degenerates into parochialism and conflict when group distinctions become explicit (e.g., Bernhard et al. 2006; Choi and Bowles 2007). This is a reason of worry, especially considering that ethnic spatial segregation is widespread in many areas of the the USA and, with some significant national differences, the EU (Musterd 2005).

Fortunately, it must be said that ethnocentrism and xenophobia are not unavoidable. Previous psychological studies showed that common categorizations, e.g., based on race, can be easily overcome by improving the information available to participants (Kurzban et al. 2001). More generally, *decategorization*—i.e., reducing the salience of group distinctions—and *recategorization*— i.e., replacing subordinate (us and them) with superordinate (we) categorizations—can significantly reduce the risks of ethnic clashes (Gaertner and Dovidio 2000; Hewstone et al. 2002). At the same time, a number of other strategies, e.g., institutional development (e.g., Ostrom 2005) and indirect reciprocity systems (e.g., Nowak and Highfield 2011), can be useful to counteract ethnocentrism without condemning societies to ethnic segregation and its risks.

This said, there is also a methodological lesson to learn from this research. While highly abstract models can offer important insights, to take them too literally can be dangerous. Although models, and especially ABMs, can play a relevant role in informing policy making, they need to be specifically designed for the task, and carefully tested against empirical data before advancing any practical suggestion (Boero and Squazzoni 2005; Smajgl and Barreteau 2014). Models are good only as far as their assumptions are tenable and all relevant elements are included. For instance, in the HG model the evolution of harming depended on the adjacent location assumption, which created a competition for space and made the use of strategies creating more room for offspring adaptive. While high competition on a specific resource may or may not be a tenable assumption depending on the case under consideration, to ignore the fact that negative actions against out-groupers often parallel cooperation with fellow group members was clearly a significant limitation of the HA model. As a result, Hammond and Axelrod's optimistic conclusions about the positive effects of ethnocentrism should be seriously put into question.

ACKNOWLEDGMENTS

The author gratefully acknowledges comments offered by Flaminio Squazzoni and by the participants in a seminar at the *Institute For Future Studies* (Stockholm, February 2014) where a previous version of the manuscript was presented.

REFERENCES

Axelrod, R., R. A. Hammond, and A. Grafen (2004). Altruism via kin selection: Strategies that rely on arbitrary tags with which they evolve. *Evolution 58*(8), 1833–1838.

Bernhard, H., U. Fischbacher, and E. Fehr (2006). Parochial altruism in humans. *Nature 442*, 912–915.

Boero, R. and F. Squazzoni (2005). Does empirical embeddedness matter? Methodological issues on agent-based models for analytical social science. *Journal of Artificial Societies and Social Simulation 8*(4), 6.

Bowles, S. (2012). Warriors, levelers, and the role of conflict in human social evolution. *Science 336*(6083), 876–879.

Brewer, M. B. (2001). Ingroup identification and intergroup conflict: When does ingroup love become outgroup hate? In R. D. Ashmore, L. Jussim, and D. Wilder (Eds.), *Social Identity, Intergroup Conflict, and Conflict Reduction*, pp. 17–41. Oxford: Oxford University Press.

Cashdan, E. (2001). Ethnocentrism and xenophobia: A cross-cultural study. *Current Anthropology 42*(5), 760–765.

Choi, J.-K. and S. Bowles (2007). The coevolution of parochial altruism and war. *Science 318*, 636–640.

Gaertner, S. L. and J. F. Dovidio (2000). *Reducing Intergroup Bias: The Common Ingroup Identity Model*. Philadelphia: Psychology Press.

Hamilton, W. D. (1964). The genetical evolution of social behavior. *Journal of Theoretical Biology 7*, 1–52.

Hammond, R. A. and R. Axelrod (2006a). Evolution of contingent altruism when cooperation is expensive. *Theoretical Population Biology 69*(3), 333–338.

Hammond, R. A. and R. Axelrod (2006b). The evolution of ethnocentrism. *Journal of Conflict Resolution 50*(6), 926–936.

Hartshorn, M., A. Kaznatcheev, and T. Shultz (2013). The evolutionary dominance of ethnocentric cooperation. *Journal of Artificial Societies and Social Simulation 16*(3), 7.

Hewstone, M., M. Rubin, and H. Willis (2002). Intergroup bias. *Annual Review of Psychology 53*(1), 575–604.

Jansson, F. (2013). Pitfalls in spatial modelling of ethnocentrism: A simulation analysis of the model of hammond and axelrod. *Journal of Artificial Societies and Social Simulation 16*(3), 2.

Keeley, L. H. (1996). *War Before Civilization: The Myth of the Peaceful Savage*. Oxford: Oxford University Press.

Kurzban, R., J. Tooby, and L. Cosmides (2001). Can race be erased? coalitional computation and social categorization. *Proceedings of the National Academy of Sciences 98*(26), 15387–15392.

Musterd, S. (2005). Social and ethnic segregation in europe: Levels, causes, and effects. *Journal of Urban Affairs 27*(3), 331–348.

Nowak, M. A., S. Bonhoeffer, and R. M. May (1994). Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Science USA 91*, 4877–4881.

Nowak, M. A. and R. Highfield (2011). *SuperCooperators: Altruism, Evolution, and Why We Need Each Other to Succeed*. New York: Free Press.

Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton: Princeton University Press.

Richerson, P. J. and R. Boyd (2001). The biology of commitment to groups: A tribal instincts hypothesis. In R. Nesse (Ed.), *Evolution and the Capacity for Commitment*, pp. 186–220. New York: Russell Sage Foundation.

Smajgl, A. and O. Barreteau (Eds.) (2014). *Empirical Agent-Based Modelling: Challenges and Solutions*. New York: Springer.

Wilensky, U. (1999). Netlogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.