

Biological Altruism in Hostile Environments

William Harms
Centre for Applied Ethics
University of British Columbia
227-6356 Agricultural Road
Vancouver, B.C. Canada V6T 1Z2
bharms@interchange.ubc.ca
619-282-5945
June 24, 1999

[12 pages, 1 table, 4 figures, approx. 3000 words]

Biological Altruism in Hostile Environments¹

Abstract: An agent-based computer model was used to investigate the effects of periodic local extinctions on the stabilization of cooperation in evolutionary prisoner's dilemma. The model allows increasing population densities which tend to undermine standard non-discriminating "solutions." The results show that for intermediate extinction rates, there is a window of opportunity for cooperation to flourish. Random extinction patterns both create environments where pure communities of cooperators survive through dissemination better than mixed and defector communities and insulate such communities from invasion. This suggests that geographical "edges of survivability" may provide the conditions under which cooperation can emerge and discriminating mechanisms evolve.

The evolution of economic altruism is one of the most vigorous areas of study at the intersection of biology, economics, and philosophy. The basic problem is easily understood. Biological organisms, be they people or paramecia, have ample opportunity to confer benefits on others at relatively low cost to themselves. If conferring such benefits becomes common, the overall productivity of the population in which it occurs is increased. Presumably, there is no advantage to refusing such benefits, but it is also the case that there is considerable advantage to pursuing a strategy according to which such benefits are accepted, but not conferred. In populations where individuals interact at random with others and when individuals play pure strategies (i.e., they always confer or do not confer benefits), the inevitable outcome is that altruistic behavior is driven to extinction. (Axelrod 1984) Nonetheless, there is ample evidence that self-sacrificing behavior is common in nature. (Sober and Wilson 1998) Consequently, individuals must not be interacting at random. The pressing question is then: what patterns of non-random interaction are responsible for the prevalence of altruistic behavior, and more theoretically, what sorts of plausible mechanisms exist that could generate the right kind of non-

¹ The research reported here was supported by Canada ????? grant # ????? and the Centre for Applied Ethics at the University of British Columbia.

random interactions required? The goal of the research reported here is to promote our understanding of altruism and guide empirical research through the identification of mechanisms that allow altruism to evolve.

The most obvious mechanisms involve individuals somehow *discriminating* as to with whom they interact (Trivers 1971), or as to whether or not they extend their altruism to particular individuals (Axelrod 1984). It has become clear that given information regarding the altruistic tendencies of others and the capacity to exploit that information, altruistic behavior can proliferate, though it is worth noting that the calculating nature of the strategies involved weakens the impression that such behavior is really altruistic. It is also widely agreed that discriminating mechanisms can be fooled, resulting in arms races of cheating and detection. (Trivers 1971)

There remains considerable interest in what sorts of mechanisms can stabilize altruism without individuals discriminating on the basis of information about those with whom they interact. In one popular model, individuals do not travel far from where they were born, which increases the probability of playing against close kin, and thus interacting with those like themselves (Hamilton 1964). Altruistic behavior benefits from the resulting correlations in such “viscous” populations. David Sloan Wilson and Elliott Sober (Sober and Wilson 1998, Wilson and Sober 1994) have militated for a rethinking of our basic picture of evolution, where the non-randomness of interaction (“population structure”) built into the representation of populations rather than into the strategies played by individuals. This sort of analysis is promoted under the heading of “group selection.”

One of the important factors that is frequently left out in the construction of these sorts of models is population density. Models of population viscosity commonly assume a constant

correlation of like with like. But when individuals tend to travel only so far from where they were born, population density increases the likelihood of interacting with other kin groups. Group selection models are in general quite sensitive to the size of forming groups (where groups are defined in terms of probability of interaction), with initially very small groups with multiple generations of isolated reproduction contributing the most to the proliferation of altruism. As population densities increase, the effective size of interacting groups increases, and the probability of multi-generation isolation decreases. The point being, that increasing population densities systematically undermine altruism in a number of models in which altruism is maintained via non-random patterns of dispersal and the spatial clustering caused by reproductive dynamics. Consequently, there is good reason to suspect that altruism might do better in hostile environments, where population densities are kept low by external forces.

The purpose of the present paper is to examine the evolution of altruism in environments where local extinction events are common using computational methods. The results are based on a non-standard agent-based computer model, where agents move between patches on a grid, forage, reproduce, and interact with other agents in the same patch. What was found is that, over a range of movement rates, extinction rates, and cost/benefit ratios, altruistic behavior not only persists but survives when selfish behavior is driven to extinction.

Simulator Design²

In the simulator used for this study, agents inhabit a ten by ten wrapped grid of patches. (Figure 1.) They move between patches with a fixed probability at each cycle (typically around one percent). Each agent has a resource level, as well as a location and a strategy. There are two strategies: always cooperate and always defect. If an agent's resource level reaches 150, the agent reproduces, breeding true. The resource of

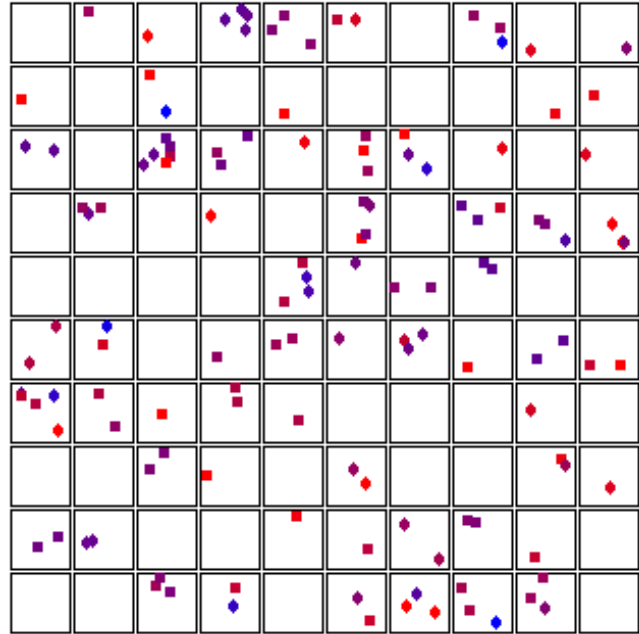


Figure 1: Agents in their patches

both parent and offspring are reset to 50, so that cost of reproduction is also 50 resource units (this value is not critical). Agents die if their resource level reaches zero, or if their patch is selected for an extinction event. Resource is accumulated in two ways.

First, agents are paired within patches at random to play one-shot prisoner's dilemma (PD) each cycle. Agents alone in patches do not play, and thus neither gain nor lose resource via play. In addition, if a patch contains an odd number of agents, one of them (again, chosen at random) will not play. Payoffs for the PD are characterized in terms of cost and benefit:

² The simulator described was written by the author in Java, using Symantec's Visual Cafe for Java. A version that runs in Java-capable browsers, as well as source code, is available at <http://eame.ethics.ubc.ca>, or from the author.

Payoffs for player 1:	Player 2: cooperate	Player 2: defect
Player 1: cooperate	benefit - cost	0 - cost
Player 1: defect	benefit	0

For the trials reported, benefit was set at 2 resource units, and costs were varied from 0.8 to 1.2 units.

Second, agents forage on a renewable resource, which regrows at a rate of one unit per patch per cycle. Agents can consume up to five units of resource per cycle, so the available resource of inhabited patches becomes quickly reduced to level of resource growth. The availability of resource for consumption obscures to a small extent the dynamics of cooperative play which is the subject of investigation here, giving advantage to the first agent into an uninhabited patch. However, this does not benefit either cooperators or defectors preferentially, and serves two other purposes. First, it allows agents alone in a patch to reproduce. Second, it increases the stability of the population at low densities, which mitigates some of the stochastic effects of the hostile environment being simulated.

As described so far, the behavior of the simulator is fairly predictable. Initial populations consisted of 50 cooperators and 50 defectors, distributed at random over the patch grid. Without exception, cooperators were driven to extinction by their costly interaction with defectors, with mean extinction times that varied from about 600 to 1000 cycles, depending on the cost of providing benefit and the movement rate.³ The critical factor in these extinctions is the increase in population densities that forces all cooperators into interaction with defectors and the attendant

³ Averages calculated for 50 trials per data point, with trials running until fixation of one of the strategies.

competition for environmental resource that might offset losses through play. Typically, cooperators were extinct by the time the global population reached 500.

Against this background, a pattern of environmental catastrophes is added. The result is an agent-based, spatially explicit version of the “metapopulation” models studied in conservation ecology (Hanski and Gilpin, 1997). In addition to the patterns of movement, interaction, and reproduction described above, each patch experienced a local extinction event with a certain fixed probability. The extinction events consisted of the elimination of all agents within the patch, regardless of their strategy. This was typically followed by a rapid replenishing of resource due to lack of consumption, immigration by one or more agents who were able to benefit from the high environmental resource level, increase in population and in mixed patches the decline of cooperators, and eventually another extinction event.

Note that if the extinction rate is high enough neither strategy can survive long. At low extinction rates the model behaves more or less as it does when there is no extinction (i.e., cooperators go quickly extinct). The phenomenon reported here happens at extinction rates intermediate between these two extremes.

A Window of Opportunity

The behavior of the model was studied over a range of values of movement rate, extinction rate, and the cost to agents of providing benefit. Generally, results were calculated over 50 trials for each set of parameter settings, though this was increased to 200 trials in some conditions (extinction rates over 1% in Figures 1 and 2) to compensate for the extreme stochasticity of the model’s behavior. Two aspects are of primary interest: the mean time to




Figure 2: Percentage of trials won by cooperators.

extinction of cooperators, and the percentage of trials in which defectors went extinct before cooperators. As noted above, at low extinction rates cooperators always went extinct first, with a mean extinction time between 600 and 1000 cycles (the left sides of Figures 2, 3, and 4). At high extinction rates, mean extinction time for cooperators converges to that for defectors (between 200 and 400 cycles) and the number of trials “won” by cooperators converges to 50%. Under such conditions, however, the winner will not last long.

Figures 2 and 3 plot percentage of trials won and mean cooperator extinction times for five movement rates over 21 extinction rates ranging from zero to two percent. Regardless of movement rate, as the rate of local extinctions increases, cooperators gradually start doing better, as evidenced by the gradual increase in mean extinction times in Figure 3. At some threshold value, cooperators begin to do dramatically better than defectors. For instance at a movement rate of .01 and an extinction rate of .011, cooperators won 86% of 200 trials. The advantage to

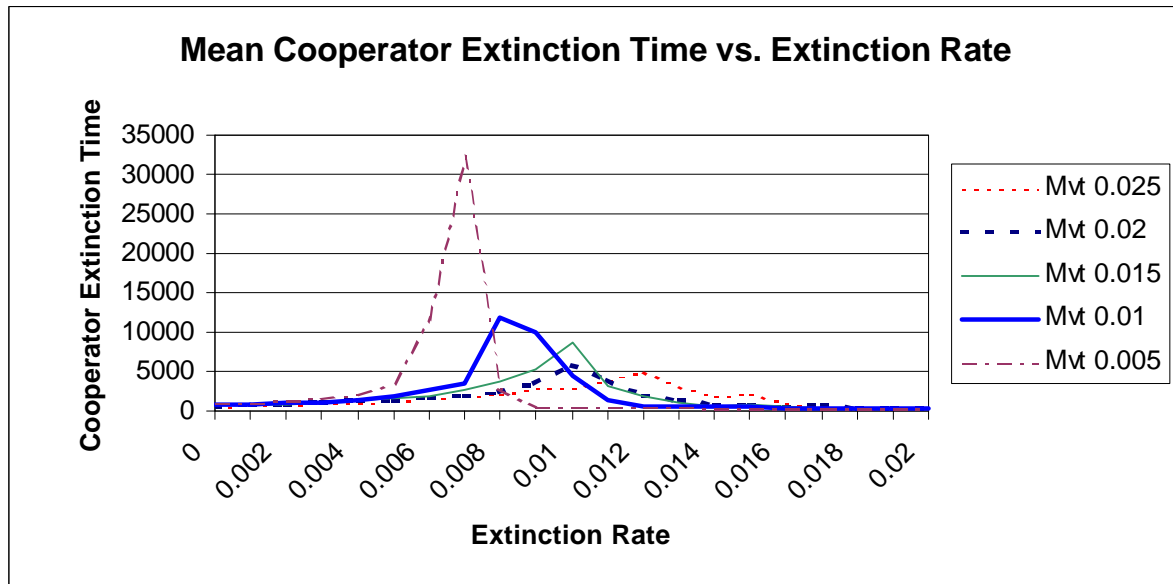


Figure 3: Cooperator extinction time.

cooperators was more pronounced for lower movement rates, and the threshold value occurred at lower extinction rates for lower movement rates. Higher movement rates result in higher dispersal and thus mixing rates, and are thus a mixed blessing for cooperators. On the one hand, greater dispersal allows both strategies to survive better by distributing the risk due to local extinctions. On the other hand, greater mixing forces cooperators into interaction with defectors with its attendant costs.

Figure 4 shows percentages of trials won by cooperators over a range of costs (combined with a benefit of 2) with a movement rate of 1%.⁴ The threshold phenomenon shows itself to be robust, though increased costs relative to benefits decrease the advantage cooperators have at intermediate extinction rates. For a cost of 0.8, and an extinction rate of 0.01, cooperators won

⁴ Note that the irregularity of the curves at high extinction rates is due to the stochasticity of the extinction pattern. This was compensated for in Figures 1 and 2 by increasing the number of trials to 200 for extinction rates over 0.01.

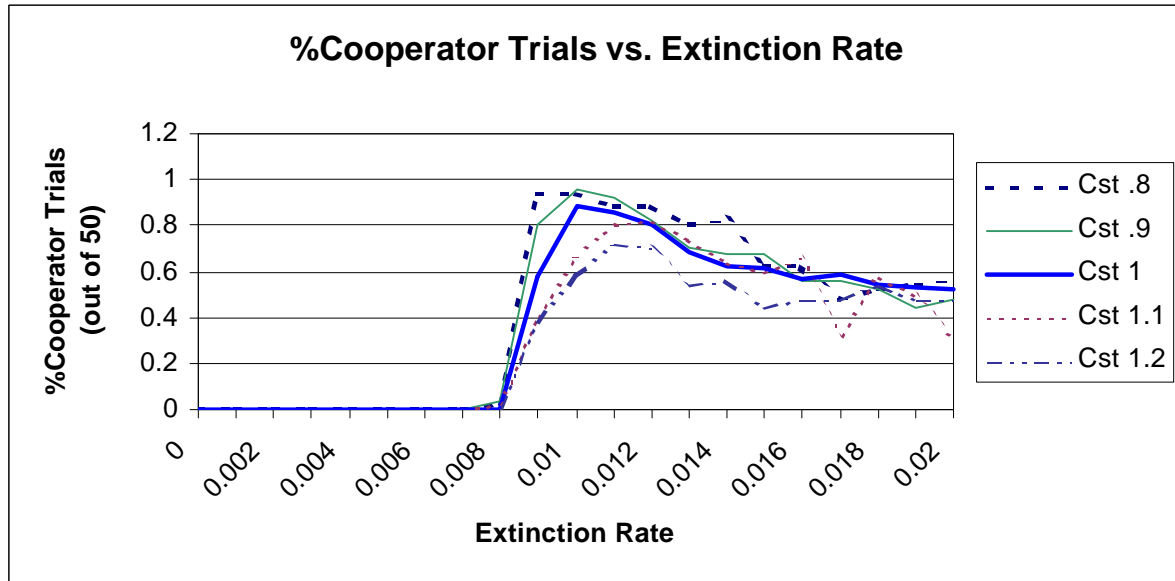


Figure 4

94% of 50 trials. Note that the range of benefit/cost ratios investigated here are rather conservative. It is not uncommon in the literature to see ratios as high as 8:1. (Cf. Hegselmann 1996.)

What This Means

In general, analytic (mathematical) methods are preferable for establishing theoretical points regarding evolutionary dynamics, and when these become intractable (as they quickly do in complex systems), computational population models offer a good substitute. Agent-based models like the one described here suffer from a high number of parameters, and a high level of stochasticity in the outcome of trials. Nonetheless, they offer a way of studying spatial effects that can at best be only approximated by population models, and the very stochasticity of the models allows us to discover effects of system stochasticity that may be “designed out” in analytic approaches. In the present case, a spatially explicit, agent-based model allows us to contrast the

way cooperators and defectors respond to stochastic features of their environment. This is not to say that aspects of the reported phenomenon are not amenable to mathematical analysis, but that agent-based models provide a sort of laboratory where phenomena can be discovered. Certainly, anything like a complete understanding will involve integration of computational and mathematical approaches.

The computational results above suggest the following interpretation: For systems relevantly similar to the model studied, as environments range from benign to hostile, there is a sort of “window of opportunity” for altruism to prosper. The effect is not critically sensitive to movement rates or costs, though these factors affect the location of the “window.”

What accounts for the effect is fairly simple. Local extinctions allow the formation of pure colonies of cooperators through recolonization. The ability to reproduce in the absence of play allows individual altruists to play only against their own offspring, or they may play against other immigrants from the same or some other pure colony. Such colonies grow faster than mixed colonies or colonies of all defectors, and since each *individual* has the same fixed probability of movement, their colonization rate is higher. This higher colonization rate allows the altruist type to avoid the inevitable extinction of its patches when defection cannot. The high rate of extinction keeps population densities low as well, around 100 agents for 100 patches in conditions in which altruists prosper. This results in a patch vacancy rate around 50%.

On reflection, it is not surprising that the greater productivity of communities of cooperators allows them to deal more effectively with hostile environments. Epstein (1998), for instance, has shown that non-discriminating cooperators can flourish in environments where defector-defector payoffs are negative, which is to say, where environments impose a fitness load

that relatively inefficient clusters of defectors cannot accommodate. What was not anticipated was that, not only are altruists able to deal with some environments better than defectors, but that spatial and stochastic aspects of those environments actually protect altruists from invasion by defectors. So, in situations where extinction and recolonization allow the formation of pure colonies (this is called the “founder effect”), and where extinction patterns maintain large number of vacant habitat niches, altruists can prosper — protected from defection by the very stochastically hostile environment that they are better equipped to deal with.

Under what conditions can we expect this phenomenon to be important? Clearly, any species or variant that survives for long must spread beyond its initial range. Its expansion will stop when it encounters waters too warm, winters too cold, or protection from predators too scarce. At these edges of survivability, altruists may have the advantage. Insofar as such edges are common, insofar as patterns of colonization allow the formation of pure colonies, altruism can gain a foothold where defection cannot, a foothold it may not be able to gain in more benign environs.

One area of further investigation is this: It is one thing to show that altruism can be stable in hostile environments, and another to show how it could have gotten established in the first place. It seems possible that edges of survivability provide conditions in which cooperative mutants can break away from non-cooperative parent populations, establishing the kinds of pure colonies of altruists that can prosper there. If this is the case, it would be worth investigating whether the partial isolation of altruist and defector communities around such edges provides conditions under which discriminating mechanisms can evolve.

Darwin (1859, Ch.3.) characterized evolution by natural selection as the “struggle for

existence.” We may struggle against competitors for scarce resources, or in more hostile environments, we may “struggle against the drought”. Too often we focus on the former, and forget the latter. In our search for fundamental principles we often average over spatial heterogeneity. In our focus on the competitive evolution of variants, we usually abstract away from absolute (and thus from small) numbers to population frequencies in “effectively infinite” populations. Much of the puzzle regarding the evolution of altruism arises from analyses that abstract away from spatiality, small numbers, open territory, and variably habitable environments. The results reported here suggest that if we look to the edges of survivability created by such factors, we may find the crucible of unselfish behavior.

References

- Axelrod, R. (1984): *The Evolution of Cooperation*, New York: Basic Books.
- Darwin, C. (1859): *On the Origin of Species by Means of Natural Selection*, ed. by E. Mayr, Cambridge Massachusetts: Harvard University Press.
- Epstein, J. M. (1998): "Zones of Cooperation in Demographic Prisoner's Dilemma," *Complexity*, 4(2), 36-48.
- Hamilton, W. D. (1964): "The Genetical Evolution of Social Behavior I & II," *Journal of Theoretical Biology*, 7, 1-52.
- Hanski, I. A. and M. e. Gilpin, eds. (1997): *Metapopulation Biology: Ecology, Genetics, and Evolution*, Academic Press.
- Hegselmann, R. (1996): "Social Dilemmas in Lineland and Flatland," in *Frontiers in Social Dilemmas Research*, ed. by Liebrand and Messick, Springer, 337-361.
- Skyrms, B. (1996): *Evolution of the Social Contract*, Cambridge University Press.
- Sober, E. and D. S. Wilson (1998): *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Cambridge, MA: Harvard.
- Trivers, R. L. (1971): "The Evolution of Reciprocal Altruism," *The Quarterly Review of Biology*, 46(3), 35-57.
- Wilson, D. S., E. Sober and (1994): "Reintroducing Group Selection to the Human Behavioral Sciences," *Behavioral and Brain Sciences*, 17, 585-654.