Altruism

A Simulated Investigation

UCLA Computer Science Honors Undergraduate Research Program
Alexander Liu and Eric Chang
Advisor: Professor Amit Sahai

Abstract

Altruism is a well documented phenomenon in species both animal and human. As defined in evolutionary biology, altruism is not just any kind behavior, but specifically is behavior that decreases the fitness of the actor while increasing the fitness of the acted-upon. Interesting facets abound when it comes to this topic, and many have been the focus of in-depth studies. By performing studies on a proprietary software platform, we hoped to gain insight on the kind of conditions under which altruism could arise and persist.

Scope

First, it must be made clear that we did not set out to answer *why* animals or people might want to perform altruistic acts, but rather *how* such acts and the genetic basis for them could be possible in the first place without being weeded out as evolutionarily disadvantageous. Furthermore, while altruism in human societies may be much more visible and relevant to someone of our species, we decided to approach the topic from a much lower level in order to avoid the numerous complicating factors. Factors such as reciprocity, reputation, cultural norms and expectations, and meme (rather than gene) transmission of behavioral traits needlessly complicate an investigation of our scope. Although we did perform several experiments in these areas on account of its undeniably fascinating implications, in the end these were deemed to be of unsatisfactory relevance to our central theme. However, as we will describe later, we would still like to see work performed in this area.

Background

Using computer simulations to study topics in evolutionary biology is not new. Thomas Ray's Tierra[1] and Michigan State University Devolab's Avida[2] are examples of existing software that utilize digital organisms. However, these are more general purpose and program-based, usually with each organism given its own discrete memory region and/or computing resources. While we did consider using a preexisting platform to begin our studies, in the end it was decided that by developing our own parameter-based program we could have a lot more specialization and control over all aspects of each experiment.

Several mathematical models are also of relevance to our investigation. Hamilton's Rule[3] is a well known statement in the form of an inequality relating the benefit of an altruistic act, the cost, and the relatedness of those involved. Similarly, the Price Equation[4] is important as a key to understanding and explaining the results of our experiments. A mathematical equation that could be used to describe the process of evolution and natural selection, the price's equation basically states the following:

The average change in characteristics between generations can be determined by the average fitness w, group specific fitness w_i, and group specific characteristics z_i.

$$\Delta z = \frac{\operatorname{cov}(w_i, z_i)}{w} + \frac{\mathrm{E}(w_i \Delta z_i)}{w}$$

If we then model the fitness of an organism in terms of its altruistic attributes:

$$w_i = \frac{n_i'}{n_i} = k - az_i + bz$$

w_i: fitness level of group i
n_i: current population level of group i
n_i': population of previous generation
a: coefficient representing loss from self altruistic actions
b: coefficient representing gain from the altruistic actions

By using the price equation, we could obtain some insights about the mathematical properties of altruism in terms of evolutionary biology:

$$\operatorname{cov}(w_i/w, z_i) = (b-a)\operatorname{var}(z_i)$$

Which essentially states that, in order for a population to have a shift in its characteristics in terms of altruism, the benefit coefficient b must be greater than the negative coefficient a.

Although it is difficult to model each specific case using the price's equation, it is still valuable for us to deduce our experimental results from a mathematical standpoint.

Method

While it is possible to perform sophisticated mathematical modeling to study such topics in evolutionary biology, we decided that we would bypass purely theoretical conjecturing by utilizing digital simulations. This would give the benefit of immediate, hands-on experimental results. The strictly numerical and parameter-based nature of our digital experiments also gave us the ability to change variables and conditions quickly and to observe the effect of these changes.

The low-level approach taken in our experiments necessitated several requirements. As fitness is defined by propagation through reproduction, and reproduction is contingent upon survival, we needed a clear measure by which organisms could survive and breed. The most clear-cut way was to quantify this was to use units of energy (the food). During each unit of time in the simulation (each iteration), the creatures would gather and expend energy for survival and/or reproduction. Next, the environment in which the organisms existed had to be defined. We decided that the environment would be compartmentalized into discrete cells (the world). During each turn, individuals could interact with the energy and with the others in the same cell, but not others. In between turns, organisms could relocate to different cells; the manner by which the relocation was determined varied between the experiments and proved to be an important parameter. In some ways this is similar to the classic Haystack Model[5], but it was when we deviated from it that we found some of our most interesting results.

Simulation Description

We set out to write a software platform that would not only implement the basic requirements described above, but also be capable of supporting more advanced features in a modular manner should they be needed. Using Java, we were able to fulfill both of these requirements. Although specific numerical parameters varied from experiment to experiment, in general a fixed world size of 1000 cells, running for at least 20000 iterations (and usually more, depending on the average time to convergence or equilibrium), with energy resources randomly distributed over the cells, was used. During each iteration, and for each cell containing food, one organism out of all those in the cell is randomly chosen to be the food's "finder." The finder can choose to keep or share this food with any others in the cell. This decision is based upon the finder's level of altruism and plays an integral part of each experiment. Of course, should an organism find itself to be the only occupant of a cell, any food in that cell is its and its alone. Food left unclaimed in empty cells roll over into the next iteration; i.e. there is no rotting of energy resources.

Reproduction occurs when an organism reaches a certain energy threshold. A child organism is split off, the parent losing an amount of energy equal to what the child starts off life with. Without mutation, the child organism shares the exact same traits as its parent. With mutation, this is almost the case, but there is a small chance that there will be a slight change in the inherited trait.

Minimal numbers of starting ancestor organisms were generated and the population was allowed to grow to around the maximum carrying capacity, which was limited by the food availability. This ensured that the final population distribution would not be skewed by the initial ratio and would instead be purely determined by how well the descendants of each type of ancestor faired. In other words, we allowed (our very simplified version of) Nature to take its course.

Finally, we named our simulation platform "eden.java", partly in whimsy with regards to the uncomplicated nature of the world and partly in reference to the Biblical command for Eden's inhabitants to "be fruitful and multiply."

General Experiment Conditions:

World Size: 1000 cells

Energy Produced Per Turn: 200 cells are chosen at random to produce 50 units of energy
Energy is cumulative, meaning that each cell could pile up the energy stock when if it is not pre-consumed by the organisms.

Organism Usage of Energy: 10 units of energy required per turn

Organism Reproduction Threshold: 200 units of energy required to reproduce

We have set up our experimental environment based upon our trials with many different variables. In general, if we do not set our world size large enough and distribute the energy uniformly, luck would become a huge factor in our final results. This is due to the fact that, in a smaller world, the organisms who by chance gain a slight upper hand can swiftly overwhelm the world with the population advantage it imposes. Upon repeated experimentation, we eventually decided upon a configuration that both gave us consistent results and took up little resources.

Typically the number of creatures each turn in our simulated environment fluctuates between 450 to 550 creatures. This matches our the theoretical threshold of 10000 (energy produced each turn) / 10 (energy consumed per turn) = ~1000 creatures support.

Changing these variables may or may not influence the results that we obtain, depending on the specific case of interest. In some cases we would discover interesting findings based on changing the test conditions. We will give specific mentioning to such cases below.

Experiment 0

The first experiment was a black-or-white case of either sharing or not sharing. In this case two ancestor organisms were generated, one with an altruism level fixed to 1.0 (the "friendly" ones) and the other to 0.0 (the "selfish" ones). The level is the fraction of food the organism shares with those with the same ancestor. As such, a friendly organism would share all of the food it found equally with any other friendlies in the same cell, while a selfish organism would keep for itself anything it found. Furthermore, there was no mutation, and organism locations were randomly regenerated after every iteration.

The result of this experiment was rather grim for the friendly individuals, as the population of completely altruistic organisms would decline to zero. This was not too surprising a result, as it was very similar to Hamilton[6] and Wilson's[7] conclusion for randomly formed groups. With this result in mind, we moved on to perform what would become the first in our main line of experiments.

Experiment 1.0

Rather than continuing with an all-or-nothing model of altruism, we decided that a more realistic approach would be to model altruism as a partial, variable behavior. This meant that depending on the level of its altruism-governing trait, an organism could share none, some, or all of the food it found. This trait ranged from 0.0 to 1.0 in increments of 0.1. For example, a creature with an altruism level of 0.4 would keep 60% of the food it found for itself, then share the other 40% with all other cell-mates, divided equally. Mutation was now an important factor. With each instance of reproduction, the child organism had a 10% chance to have its altruism level increase or decrease by 0.1 from that of its parent.

The results of this experiment were as such: with enough food to support a population of about 1000 individuals, on average about 600 would have an altruism level of zero, 300 with 0.1, and so on as shown in the graph below.

On the whole this seemed to support the idea that perhaps a slight degree of altruism can be supported. At the very least, slight altruism to even complete strangers would not drastically reduce the fitness of the organism. However, this conclusion would refined and further developed with some of our later experiments.

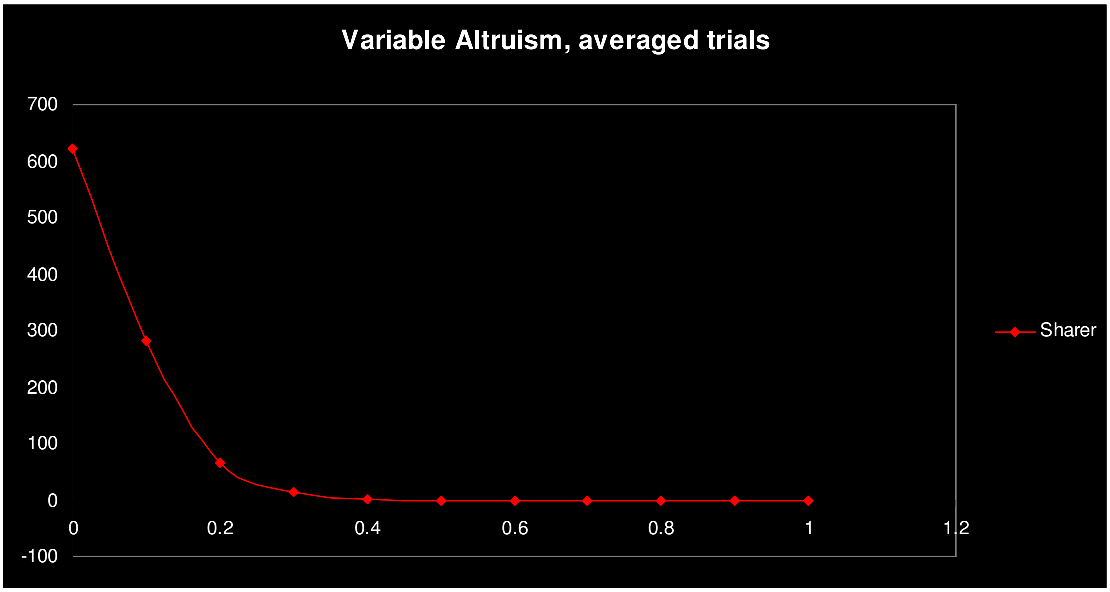

Experiment 1.5: Modifications

We made a series of modifications to the previous experiment in order to determine the difference between altruistic behavior towards random strangers and towards kin. The same variable now could govern one of or both of the following behaviors: the same sharing behavior as described previously, or the proportion of food shared by a parent organism with its child organism. We called this the "breeder" behavior. The parent would keep half its energy, and distribute the rest according to its level of altruism to its child. A parent at 1.0 would give 50% of its total food to its child when reproducing, one at 0.5 would give 25%, and one at 0.0 would give nothing (ensuring the almost immediate death of its child). The results of these modifications are shown in the graph below.

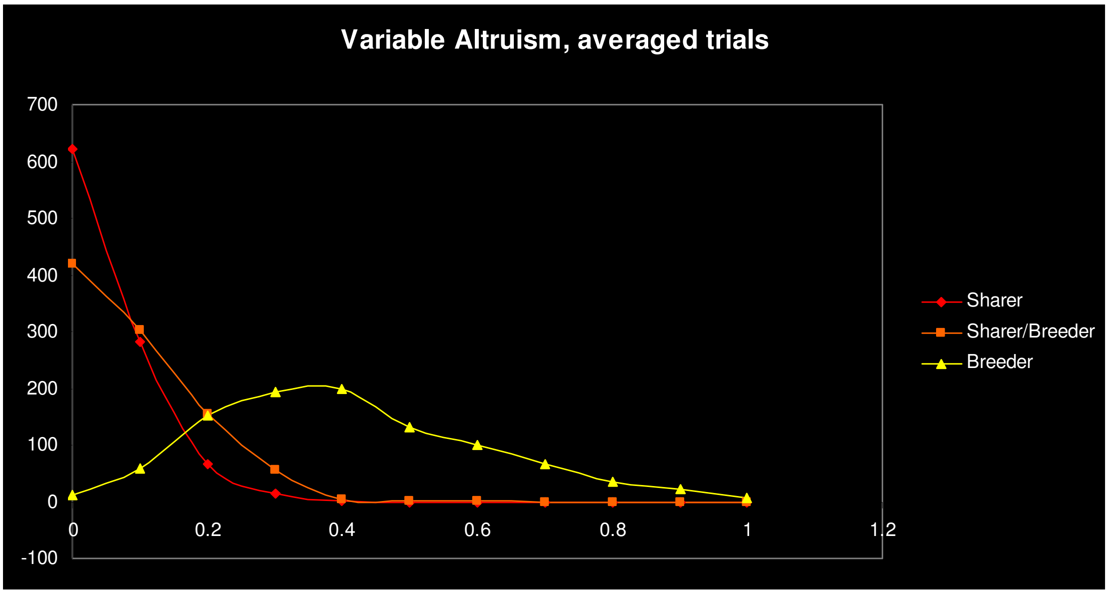

As we can see that graph, the beneficial effect of being more altruistic towards one's own children is immediate obvious. Organisms which were completely selfish with strangers but behaved nicely towards their progeny tended to end up with a higher level of "niceness" than organisms which behaved identically to their progeny and similarly towards strangers.

Experiment 2.0 and 2.5

In order to better simulate the effect that real genes have upon behavior in real life, we next implemented the level of altruism not as a function of a simple sliding variable, but of the sum of a set of on-off "genes". An organism would start with a gene-set of ten genes, each in either an "on" or "off" state. The creature's level of altruism would be equal to the number of "on" genes it possessed. For example, an individual with six "on" genes and four "off" genes would have an altruism level of 0.6. Similar to the previous experiments, there was a slight chance of mutation. During reproduction, the child organism would have a 10% chance of one of its genes (randomly determined) change state.

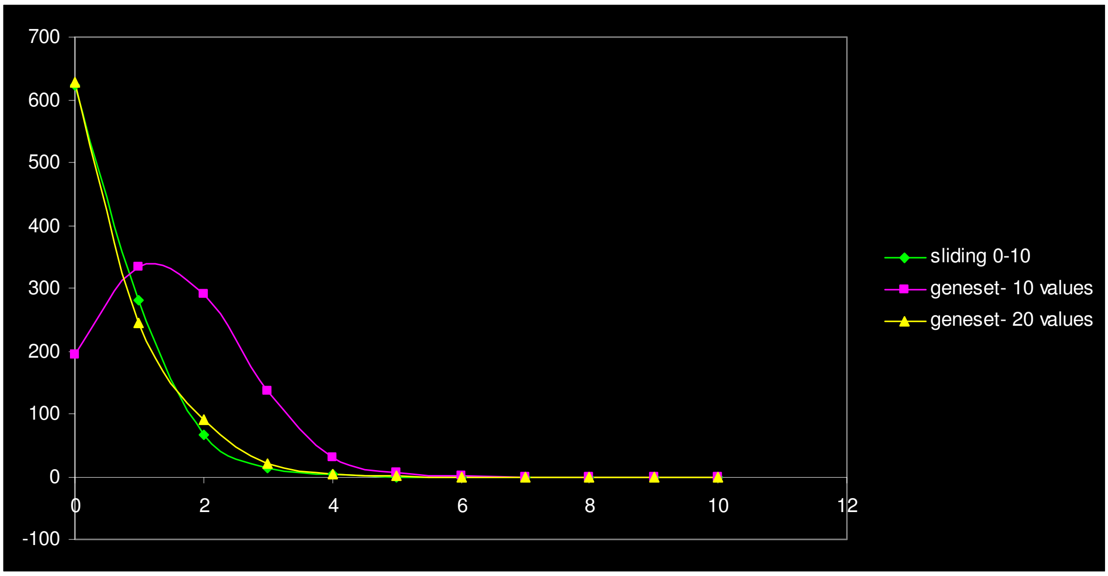

The results of this experiment are shown above. In green is the result of Experiment 1.0 shown for comparison. In purple is this experiment carried out with a gene-set size of ten.

We determined that with ten genes, the results were colored by the simple mathematical fact that the descendants of an organism at one end of the spectrum would be predisposed to migrate towards the center through mutation. For example, if mutation was to occur during the reproduction of the child of a parent with nine of the ten genes turned off, chances would be that the mutation would turn one gene on. As such, we carried out a modification of this experiment. In the graph above, the yellow line shows the same experiment modified to have a gene-set of twenty. The first and last five genes to be turned on would have no additional effect. In effect, only the middle ten genes determined the expressed behavior. An organism with eight genes turned on would have a level of altruism equal to 0.3, one with fourteen would be at 0.9, and one with eighteen would be at 1.0.

After the modification, the result of the experiment hewed very closely to Experiment 1.0. Again, this seemed to show that a small degree of altruism could be supported in the environment. Exactly how much and whether there was an absolute advantage over simply being selfish remained to be seen in the next few experiments.

Experiment 3.0

The results we have obtained so far showed that there is always a slight population of altruistic organisms in the equilibrium for each specific case. We postulated that this is simple a slight aberration due to our mathematical model: because of the pre-set mutation rate, there will always be some organisms that end up being slightly altruistic. We thus designed our next set of experiments to limit the factor of random mutation rates. We increased the range of possible altruistic characteristic values from 1~ 10 to 1~100. Thus, an aberration in mutation would only induce a slight insignificant shift in the altruistic values of the population, instead of the very significant 10%~20% we have observed in previous cases.

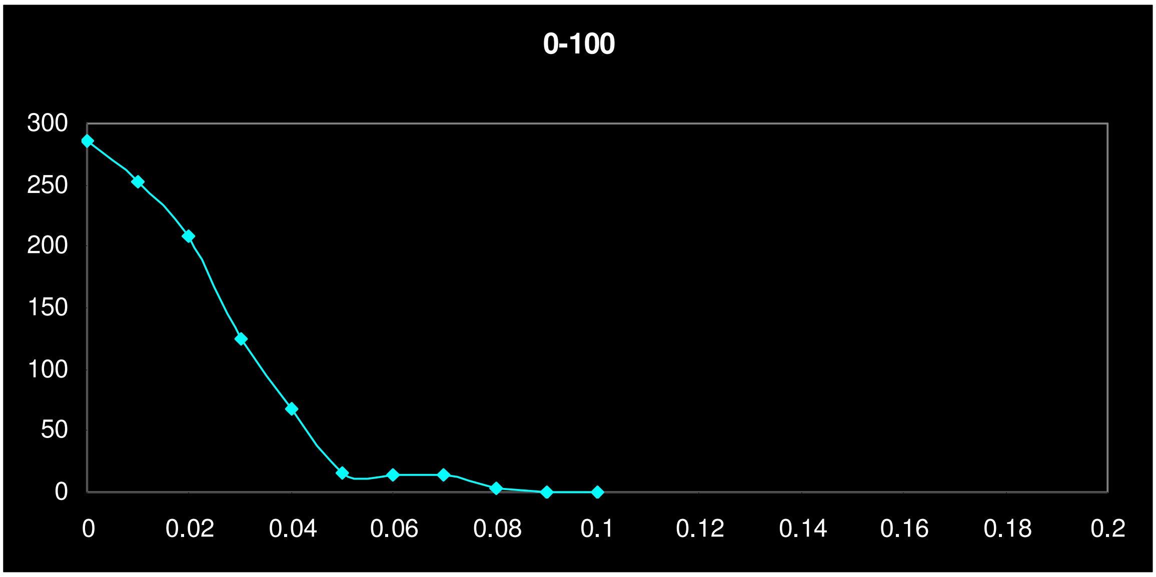

This time, the largest altruistic value that an organism can reach is capped out at 0.5%. Despite this, one can still draw some meaning conclusions from this observation: altruist organisms can exist even when they obviously decrease the fitness of the altruists. The decrease in fitness may be so insignificant that it would not hurt the organism to share a little.

Experiment 4.0 and 4.5

It is not surprising that, in all the previous cases, there is little altruism to be observed. Intuitively, there is simply no benefit for being purely altruistic: the selfish organism would simply benefit from the altruistic actions and enjoy a larger fitness while the altruistic organisms would suffer from reduced fitness. Thus, we decided to implement a conditional altruism for our organisms. Namely, we set a restriction for which the altruistic organisms are to share energy with each other: the organism will only share its findings with other organisms that have an altruistic value that is within the range of +- 0.5 with their altruistic value.

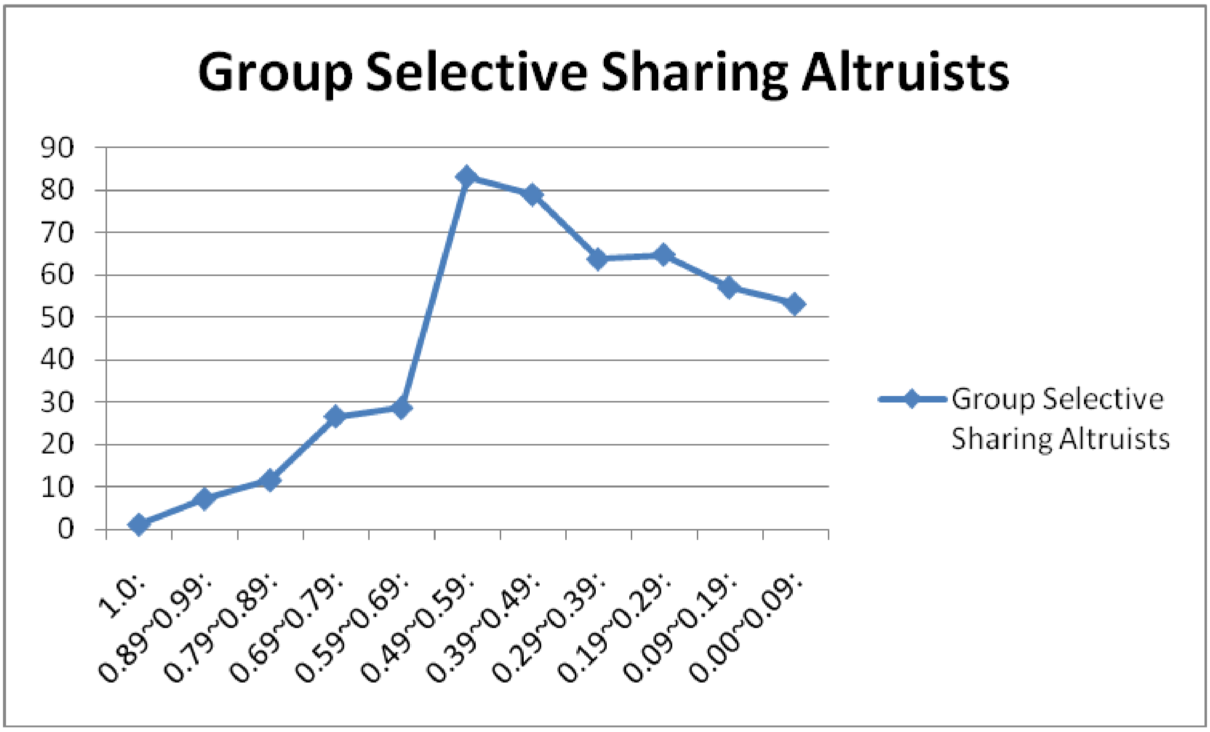

The results we obtained did indeed stray away from the previous cases that we have tested. As it can be observed from the graph, the overall distribution of the altruistic characteristics can be described as a skewed normal distribution. In the context of the Price's equation, altruistic traits that arise in a population strongly depend on how much benefit is gained from group based altruism.

The experiment would lead us to reinvestigate our results from experiment 0. Experiment 0 pitted completely altruistic organisms against completely selfish organisms, but the conditions for altruistic actions are very similar to the case for Experiment 4, where they only share with closely related peers. One major distinction between the two experiments, however, is that while the movements of organisms in experiment 0 are completely random and bounded only by the size of the world, we restrict the movements of organisms to uniformly random from -5 to +5 change in position.

Snapshot of the population distribution at equilibrium

Location	# of Altruist	# of Selfist
0	91	0
1	98	0
2	97	0
3	103	8
4	56	34
5	31	72
6	0	94
7	0	105
8	0	92
9	0	105

*the world is divided into 10 regions based on locality.

Snapshot of an invasion taking place

Location	# of Altruist	# of Selfist
0	102	0
1	95	0
2	107	0
3	99	0
4	105	0
5	75	25
6	74	19
7	45	54
8	9	96
9	0	86

We were able to draw several interesting observations from the experiment. For one, the altruistic organisms no longer get dominated by their selfish counter parts. Instead, it is now the altruistic organisms that consistently takes over the simulated environment. In most test cases, the organisms would reach equilibrium at a 1:1 ratio. A “tug of war” would then follow, in which the population of both sides would fluctuate up and down. We also observed a very interesting and unexpected behavior: the two types of organism would separate in distinctive geological groups, each taking one half of the world and control their “territory.” An invasion would then take place, in which the population of one side would begin to show increase in the “no man’s land” between the two territories. What we have observed is that selfish organisms generally have a difficult time holding off an invasion, they die out rather quickly once the altruists gain a sufficient population upper hand. The altruists, however, are much more robust in their resistance to invasions, and generally have the ability to rebound back to the equilibrium. Thus it could be said that altruism may increase the fitness of an individual only if the individual is working within the context of a “group” competing against other groups.

References

1. Tierra Home Page. Ray, Tom. 2006. < http://life.ou.edu/tierra/>

2. Lenski, Ofria, Pennock, Adami, Avida, 2003. The evolutionary origin of complex features. Nature Vol 423 8 May 2003 pp.139-143.

3. Haldane, JBS. 1955. Population Genetics. *New Biology* **18**:34-51

4. Gardner, A. (2008). "The Price equation". *Curr. Biol.* **18**: R198-R202.

5. Wilson, D 1987. Altruism in Mendelian Populations Derived from Sibling Groups: The Haystack Model Revisited. *Evolution*, Vol. 41, No. 5 (Sep., 1987), pp. 1059-1070

6. Hamilton, W.D., 1975. Innate social aptitudes of man: an approach from evolutionary genetics. In: Fox, R. (Ed.), Biosocial Anthropology. Wiley, New York, pp. 133–155.

7. Wilson, D.S., 1975. A theory of groupselection. Proc. Natl Acad. Sci. 72, 143–146.