

Population Viscosity and the Evolution of Altruism

Joshua Mitteldorf* and David Sloan Wilson**

*Department of Biology, Leidy Laboratory
University of Pennsylvania
Philadelphia PA 19104
josh@forum.swarthmore.edu

** Department of Biology
SUNY
Box 6000, Binghamton, NY 13902-6000
dwilson@binghamton.edu

Abstract

The term population viscosity means limited dispersal, which increases the genetic relatedness of neighbors. This both supports the evolution of altruism by focusing the altruists' gifts on relatives of the altruist, and also limits the extent to which altruism may emerge by exposing clusters of altruists to stiffer local competition. Previous analyses have emphasized the way in which these two effects can cancel, limiting the viability of altruism. These papers were based on models in which overall population density was fixed. We present here a class of models in which population density is permitted to fluctuate, so that patches of altruists are supported at a higher density than patches of non-altruists. Under these conditions, population viscosity can support the selection of both weak and strong altruism.

Introduction

The emergence of altruism is a fundamental problem in evolutionary biology: How can nature select a gene that promotes the fitness of others, especially when it is at the expense of the bearer of the gene itself? At a minimum, the benefits of the altruism must fall disproportionately on other altruists. This may be by virtue of intent, as in reciprocity models that derive from game theory; or it may result from explicit relationship of the altruist to his beneficiaries, as in kin selection models; or it may be facilitated by social group structures that support the altruistic behavior and evolve in parallel. Each of these ideas has been modeled and studied in recent years, with results that are encouraging. In the present study, we seek a paradigm that is both simple and generally pertinent, in which altruism may be selected on the basis of geography alone. Population viscosity simply means the tendency of offspring to disperse slowly from their site of origin, so that average relatedness correlates with geographic proximity. This affects the prospects for selection of altruism in two ways: If the altruist's gift is dispensed locally, other altruists are more likely to benefit; this promotes the selective prospects for the altruistic trait. But if competition

also takes place principally among neighbors, then viscosity slows the adoption of any beneficial trait, altruism included, by exposing the advantaged population to stiffer competition.

The present work grows out of a model published by Wilson, Pollock and Dugatkin (1992). WPD reported that these two effects tend to cancel, and that population viscosity models were not promising for supporting the evolution of strong altruism. But their conclusion depends critically on a non-essential feature of the model: that overall population density was fixed, intensifying local competition. It is easy to think of examples of altruistic traits that increase the viability of clusters of altruists in such a way that they are supported at a higher density than similar clusters of non-altruists. We have found that by simply permitting population density to vary, we arrive at a conclusion far more promising than WPD. The results depend on a variable overall population density. A full account of the present work appears in *J. Theor. Biol.* for June of this year (Mitteldorf & Wilson, 2000).

The model

Here's how the model works: Every site on a Cartesian grid may be occupied by either an altruist (A) or a non-altruist (S) or it may be empty (V). Grids are two dimensional, 200x200 in most runs. The occupants reproduce asexually in fixed, non-overlapping generations. Both competition for the grid site and dispersal of altruistic benefit are local, limited to a von Neumann neighborhood (or VNN, five sites in the shape of a "+").

At each time step, a lottery is held to determine whose offspring becomes the next occupant of each grid site. Participants in the lottery are occupants of the site itself and its four lattice neighbors. The lottery is weighted by the fitness of each of the five participants, computed as follows: The base fitness of a non-altruist is 1, while that of

an altruist is $1-c$, where c is the cost to the self of altruism. To the base fitness is added an altruistic benefit b , scaled by the number of altruists in the VNN centered on each individual.

$$W_A = \text{fitness of altruist} = 1 - c + b(N_A/5)$$

$$W_S = \text{fitness of non-altruist} = 1 + b(N_A/5)$$

where $N_A/5$ is the average proportion of altruistic neighbors (the VNN includes the self). Note that five overlapping VNNs are used to compute contributions to the fitnesses of the five participants in the lottery that determines which will seed the central site in the new generation.

The model thus constituted makes no provision for empty sites, and the total population density is implicitly fixed. Fitness in this version of the model determines the relative advantage of each variety in competition with the other, but in absolute terms a fitter variety cannot populate the grid at a higher density. Just this model was analyzed by WPD in 1992, and it affords a limited hospitality to altruism. For parameters b and c that are small compared to unity, the model permits altruism to emerge whenever $b > 5c$. This happens to be the boundary between "weak" and "strong" altruism, because $1/5$ of the local benefit b contributed by each altruist is enjoyed by the altruist itself. (We say "happens to" because with minor geometric variation of the model, there is a similar minimum b in relation to c , but it does not coincide with the weak/strong boundary.)

Last year, Wilson and I extended this model to allow for vacant sites, and therefore variable population densities. We did this in three separate ways, and all had the effect of making the model much more hospitable to altruism. The rules governing the lottery need to be generalized to allow for vacant sites (V's) as well as A's and S's. The simplest rule would be to assign a constant "fitness" η to the void; each V participating in a lottery would receive η tickets, independent of its surrounding neighborhood. The language describing V's as a separate species with its own fitness is a convenient mathematical fiction, akin to the symmetric treatment of negative electrons and positive holes in a solid state physicist's equations for a semiconductor.

One more ingredient needs to be added in order to have a non-trivial model with vacancies. If the fitness of the void η were set >1 , then all populations would disappear; if η were set <1 , then the grid would fill to saturation (voids would "become extinct"). (The borderline level $\eta=1$ is unstable to either of these extremes.) One way to follow evolution for an extended period on an unsaturated grid with $\eta < 1$ is to introduce periodic disturbance events. On a fixed schedule, a percentage of all grid sites is vacated, as if by a disease or natural catastrophe, or a seasonal change that kills off a

high proportion of individuals in winter and permits regrowth in summer. The geometry of the disturbance may be that a fixed proportion of sites is chosen at random for evacuation ("uniform culling"); or at the opposite extreme, all the sites in one area of the grid may be vacated, while the sites outside of this region remain untouched ("culling in a compact swath"). Both these cases have been explored, and results are reported below.

Periodic uniform culling

For the uniform culling case, we choose parameters such that a high proportion of the population is culled, denuding the grid except for isolated individuals. This maximizes the founder effect, as regeneration takes place in patches that are likely to be pure A or pure S, because they are descended from a small number of survivors. As patches begin to regenerate, they expand into surrounding voids, and direct competition is delayed for a time. In the initial phase of regrowth, it is the competition A vs V and S vs V that dominates; only later do A's and S's compete head-to-head. But A communities, by our assumptions, can grow faster than S communities; hence the A's may compile a substantial numerical advantage before direct encounters (A upon S) become commonplace.

Population culling in a compact swath

The other possibility we have explored is culling in a compact swath. A square centered on a random point and sized to include half the grid's total area is completely denuded in each disturbance event. Regeneration takes place from the edges, and regions of the boundary that, by chance, are dominated by A's expand most rapidly into the void. Although the founder effect is evoked less explicitly in this variation than in the uniform culling case above, culling in a swath nevertheless proves to be a surprisingly efficient mechanism for segregating the population into patches of pure A and pure S.

Elastic population densities in steady state

In a third variation, we introduce an extra parameter that permits voids to coexist with the other populations in steady state. The parameter ξ represents a constant extra chance for the void in each 5-way lottery, an acknowledgment that there is always a finite chance that seeding will fail and a site will become vacant, even if its neighborhood is saturated in the parent generation. Specifically, the rules of the lottery are modified so that in addition to the 5 competitors comprising the A's, S's and V's in a given neighborhood, there is always an extra measure of ξ lottery tickets assigned to V. Now V can never become extinct because each lottery has a guaranteed minimum probability of generating a V. But if η is small enough, then A's and S's will still have an advantage in each lottery sufficient to insure that they do not succumb to V in competition. With proper choice of η and ξ a stable steady state results. This third model supports the emergence of

altruism in appropriate parameter ranges, and favors quasi-periodic population oscillations for other ranges. Overall, we have found it to be a rich lode of subtle and unexpected phenomena.

Results

Much of the interesting detail relevant to the success of altruism is contained in the part of parameter space close to $b=5c$, the boundary that divides weak from strong altruism. In order to explore the region of parameter space near to $b=5c$ for further study, we introduce the parameter γ , defined as the quotient of the total benefit to others by the net cost to the self.

$$g = \frac{4b/5}{c - b/5} = \frac{4b}{5c - b}$$

In the results tabulated below, we focus on "stalemate γ " as a measure of each model's hospitability to altruism. Stalemate γ is located as the end result of multiple trials, seeking that combination of b and c for which altruists and non-altruists have equal prospects.

Replication of WPD model results

Our first task was to replicate the results of the 1992 model. This we did for small b and c . For larger b and c , we discovered a new behavior: the viability of altruism increases slowly as b and c are raised in tandem, maintaining a fixed ratio. Stalemate γ is infinite for small c , meaning that only weak altruism is viable; but as c rises to unity, we find γ is about 91, and at $c=10$, γ is about 33.

Periodic uniform culling

In these runs, a high percentage of the population was culled at random on a regular schedule, such that much of the action consisted of regrowth from isolated founders into empty surroundings. Much of the competition is not directly A against S at their common border, but rather is a race for free expansion into unpopulated regions. Higher values of η imply slower growth rates for both varieties, but as η approaches 1, altruists are affected relatively less; this is because at $\eta=1$, non-altruist populations do not grow at all, but altruist populations may still grow for somewhat higher values (depending on parameters b and c). Therefore, higher η and deeper, more frequent culling are the conditions more favorable to the emergence of altruism, consistent with our model results. We found values of stalemate γ ranging down to 3.6 for 95% culling every 100 generations.

Population culling in a compact swath

In this model variation, we specify that cullings cut a square swath across the grid in which all occupants are removed. This scheme was found to be an efficient segregating force,

creating within a few culling cycles very clearly-defined regions filled densely with homogeneous populations of type A or S. Some competition takes place at boundaries between A and S regions, but what is more important turns out to be the rate at which each group regrows into the voided swath. As in Section 3.2, the frequency of culling determines the balance between these two forces.

Because of the segregation, these runs were surprisingly hospitable to the evolution of altruism. Compared with uniform culling, we find a much lower stalemate γ for the same severity of culling; in fact, experiments with a compact culling factor of 50% are even more hospitable to altruism than are uniform culling models in which only 5% are left standing. Growth into a vacated region from the edges takes much longer than filling in an equivalent area of smaller, distributed voids, and hence is a more sensitive test of a population's ability to expand.

Two trends were identified: First, increasing η with everything else held constant creates an environment more hospitable to altruism because the growth rate of non-altruist patches is reduced while the growth rate for altruist patches is less affected. Second, increasing the frequency of culling also enhances the viability of altruism, by increasing the fraction of the time that growth is taking place unimpeded into void regimes.

Another clear trend noted in our results is that the outcome of individual runs is less predictable in this section (with culling across a swath). In other words, the range of γ values for which A or S may randomly prevail in a given run is largest with this paradigm. We speculate that the correlation between maximum effectiveness of group selection and minimum predictability is a broad trend. The problem of group selection may be stated: how can individual selection, which is quicker and more efficient, be forestalled long enough for intergroup differences to show their effect? This formulation suggests that any stochastic effect, making the short-term outcome less certain, is likely to decrease the importance of individual selection relative to group selection.

Elastic population densities in steady state

Models in this section allow for variable population densities, with patches of altruists supported at a higher density than patches of non-altruists, but the population is not continually expanding as in the above variation. Rules of the lottery have been modified as described in Section 2.3 above so that voids (V) co-exist with A's and S's in steady state. In the models with culling, much of the competition takes place for expansion rate into large voided areas; but here, in contrast, all of the competition is head-to-head, A vs S, albeit in the presence of scattered voids. How can A outcompete S in these circumstances? At the battle front between colonies of A and of S, the A's are arrayed more densely, because their cooperative benefit

permits higher steady-state population densities in areas dominated by A. The A's tend to have a local numerical advantage at these boundaries, which can offset the fitness advantage of S, which is always present.

The major trend we identified in these runs is that higher grid vacancy implies a more hospitable environment for altruism: staminate γ decreases roughly exponentially with increasing grid vacancy. The reason is that density elasticity is the driving force behind the success of altruism, and the dynamic range of population density is limited by the assumption of one or zero occupants per site. Maximum population density is 100%; minimum population density is about 30% because when parameters are adjusted for a density lower than about 30% the population is too fragmented to survive. This suggests that we have only begun to explore the population elasticity variable, with a dynamic range of just a factor 3. The minimum $\gamma=4.7$ that we observe may well be surpassed in model variations that allow for a population variable at each grid site.

Discussion of results

The 1992 conclusion of WPD, using a model with fixed total population density, was considered discouraging for general explanations of altruism with population viscosity. Our motivation for modifying the WPD model was the insight that variable population density might be a crucial factor for the viability of altruism. We began with models in which groups of altruists, though unable to compete head-to-head against non-altruists, were nevertheless able to grow into a vacant habitat at a faster rate. We found that models in which the populations were periodically culled, creating empty space to be repopulated, constituted a friendly environment for the emergence of altruism. We moved from there to model variations in which dispersed vacancies were built into the population, in such a way that clusters of altruists could exist in steady state at densities up to 3 times the corresponding population density for clusters of non-altruists. We saw that altruists were able to succeed in this model through a mechanism that was somewhat less transparent than that of free expansion: First, viscosity supports division of the lattice into patches dominated by one or the other variety. The fitness advantage which altruists confer upon their neighborhood permits those patches dominated by altruists to establish denser populations. In the competition that takes place at patch boundaries, the greater density of altruists permits them to counteract the fitness advantage of non-altruists in close proximity, and in some cases to prevail.

Population viscosity models, because they assume only a two-dimensional geography and limited dispersal speed, are a formalism of great generality. We have seen that weak altruism may be supported in these models quite generally, and that the sorts of strong altruism that can be supported are such that the altruistic trait must contribute to a higher

population density in patches of altruists than exists in patches of non-altruists.

The academic community of evolutionists debated the character of natural selection in the 1960's and 70's, arriving at a skepticism concerning group selection that continues to inform theoretical work through the present. The debate took place at a time when analytic models in the tradition pioneered by R. A. Fisher were the state of the art: Population growth was treated differentially, using continuous functions to approximate their discrete analogs. Large population size was implicit to this framework, and random mating was frequently invoked. In the ensuing decades, computers have become ubiquitous and large-scale modeling has become practical and generally available. The words "chaos" and "complexity" have emerged into common parlance. With simple computer models, we find a basis for a more generous view for the role of group selection in nature. One of our conclusions is that stochasticity itself is a factor favoring the selection of altruism. Perhaps the present availability of computer models is sufficient reason to re-open the group selection debate and re-examine its essential conclusions.

References

- Mitteldorf, J.J. and Wilson, D.S. 2000. Population Viscosity and the Evolution of Altruism. *J. Theor. Biol.* 204:481-496.
- Wilson, D.S. Pollock, G.B. & Dugatkin, L.A. 1992 Can altruism evolve in purely viscous populations? *Evol. Ecology* 6:331-341.