

A Mechanism for the Evolution of Altruism Among Non-kin: Positive Assortment Through Environmental Feedback

John W. Pepper
Barbara B. Smuts

SFI WORKING PAPER: 2000-12-065

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

A mechanism for the evolution of altruism among non-kin: Positive assortment through environmental feedback

John W. Pepper
Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501
jpepper@santafe.edu

Barbara B. Smuts
Department of Psychology
The University of Michigan
Ann Arbor MI 48109
bsmuts@umich.edu

Abstract

The evolution of cooperation often requires genetic similarity among interactors. For populations that are divided into groups such that a given trait affects all group members, this requires positive assortment of individuals into groups, i.e., that individuals are more similar to other group members than to the population at large. Several authors have claimed that mechanisms other than kinship could produce genetic similarity within groups, but this claim has not been generally accepted. Here we describe such a mechanism. The process of “environmental feedback” requires only that the cooperative trait affects the quality of the local environment, and that organisms are more likely to leave low quality than high quality environments. We illustrate this dynamic using an agent-based model of feeding restraint. The mechanism of environmental feedback appears to be a general one that could potentially play a role in the evolution of many kinds of altruism in nature.

Key words: group composition, inclusive fitness, kin selection, group selection, multilevel selection, cooperation

Introduction

Cooperation has been a central problem in evolutionary biology for several decades. The paradigmatic example is the evolution of altruistic traits, which impose a fitness cost on the “actor”, or the individual expressing the trait, while providing a fitness benefit to one or more “recipients”.

One explanation for how altruism can evolve assumes that individuals strategically modify their behavior in response to others’ past behavior. This forms the basis for the theory of reciprocal altruism (Trivers 1971), and a large literature on evolutionary game theory (Axelrod & Hamilton 1981, Axelrod 1984). Most traits of organisms may not involve this type of cognitive flexibility however, and here we address the evolution of altruism in fixed or genetically determined traits. In this context, the most common explanation for the evolution of altruism is the theory of inclusive fitness (Hamilton 1964). The alternative framework of multilevel selection theory (Price 1972), though different in approach, has proven to be mathematically equivalent (Hamilton 1975, Queller 1992).

Under both the inclusive fitness and multilevel selection frameworks, whether altruism evolves depends critically on genetic similarity between actors and recipients. When populations are divided into groups such that social traits affect all group members, this genetic similarity depends on the pattern of assortment, or how individuals are distributed among groups with respect to the trait in question. When individuals are more similar on average to other members of their group than to members of the population at large, assortment is said to be positive. When individuals are less similar to members of their group than to the population at large, assortment is negative.

The central role of assortment can be understood through either the inclusive fitness or the multilevel selection framework. In terms of inclusive fitness, positive assortment increases the relatedness coefficient between actors and their recipients (Hamilton 1975; Pepper 2000). In terms of multilevel selection, positive assortment increases genetic variance between groups and reduces it within groups, thus strengthening between-group selection at the expense of within-group selection (Price 1972; Hamilton 1975; Wilson 1977). Regardless of which framework is used to understand this effect, positive and negative assortment have strong positive and negative effects, respectively, on the evolution of any altruistic trait.

Traits that benefit only other individuals and not the actor are “strongly altruistic” (Wilson 1979), i.e., they impose a net fitness cost on the actor. Such traits cannot be selected for without positive assortment. In contrast, traits benefiting a group that includes the actor may be only “weakly altruistic”, meaning that they impose a fitness cost relative to the rest of the group, but not relative to the global population. Such traits can spread without positive assortment, but it greatly facilitates their evolution as well (Wilson 1990, Pepper 2000).

Given the crucial role of assortment, the question of how it can arise becomes central. In his early publications Hamilton focused solely on genetic similarity through common descent (Hamilton 1964). For this reason Maynard Smith (1964) renamed inclusive fitness as “kin selection”. But according to Hamilton, “Kinship should be considered just one way of getting positive regression of genotype... Thus the inclusive fitness concept is more general than ‘kin selection’.” (Hamilton 1975).

Hamilton did not discuss in any detail what other mechanisms could produce this “positive regression of genotype”, or positive assortment. He suggested, however, that either pleiotropic effects on habitat preference or the ability to recognize and associate with other altruists might suffice. Wilson (1977) explored the idea of non-random assortment through environmental interactions in more detail. However, no mechanism has been proposed to prevent the evolution of “cheaters” that combine selfish behavior with the traits leading to association

with altruists. Because of such considerations, Maynard Smith (1976) argued that, “these other reasons for genetic similarity between neighbors seem likely to be unimportant”. Grafen (1984) concurred, stating that, “The only plausible cause of genetic similarity among group members is common ancestry.” This remains the standard view today, and most reviews and textbooks do not mention even the possibility of genetic similarity without common descent.

More recently, several models have pursued the idea that altruists could actively choose to interact with each other, either in general terms (Eshel & Cavalli-Sforza 1982, Wilson & Dugatkin 1997), or through specific mechanisms such as reputation (Nowak & Sigmund 1998), or direct observation and memory of others’ behavior (Cox et al. 1999). These mechanisms require substantial cognitive abilities however, and therefore probably apply to few non-human species.

Here we propose a source of positive assortment that does not require either common descent or special cognitive abilities. This mechanism, which we refer to as “environmental feedback”, requires only that the trait in question alters the local environment in some way, and that the organisms’ movement patterns respond to this change.

The specific altruistic trait we study here is feeding restraint, also called prudent predation (for review see Hart et al. 1991). When a food source has a non-linear (e.g., logistic) growth function, lower feeding efficiency can increase the rate of food production. Feeding restraint is thus an altruistic trait, creating both an individual cost in terms of lowered food intake relative to competitors, and a group benefit in terms of increased food supply to all individuals in the area (Gilpin 1975).

In an earlier paper, we showed that in patchy environments feeding restraint could spread under a fairly wide range of conditions, even without kin selection (Pepper & Smuts 2000). One reason is that the recipients of this trait’s benefits include the actor, which can lead to positive relatedness coefficients between actors and recipients even under random assortment (Pepper 2000). A hypothesized second reason is that the biologically plausible migration patterns in this model produced positive assortment without common descent. That hypothesis is the subject of the current study.

Methods

This study was based on an agent-based computer model described previously (Pepper & Smuts 2000). The program is written in Objective C using the Swarm library for agent-based modeling (Minar et al. 1996), and is available on request.

The model world and agents

The model consisted of a two dimensional grid, wrapped in both axes to avoid edge effects, containing two kinds of agents: plants and foragers. Plants had a fixed location in a given grid cell, and varied only in their size, which represented the amount of food energy available to foragers. At the start of a run each plant’s initial size was set to a uniform random number between zero and a fixed maximum. At each time step this energy level increased according to the logistic growth curve:

$$\Delta S = RS(K - S) / K$$

where ΔS = growth per time step, S = current size, R = logistic rate of growth, and K = maximum size.

Foragers moved according to the following rules: They examined their current and eight adjacent cells, and from those not occupied by another forager, chose the cell containing the plant

with the most energy (with ties broken randomly). If the chosen cell would yield enough food to meet their metabolic cost for one time step they moved there. (This fixed metabolic cost per time step was the same for all foragers.) If not, they moved instead to a randomly chosen adjacent cell not occupied by another forager. This movement rule led to the emigration of foragers from depleted patches. This simulated the behavior of individuals exploiting a local food source while it lasts, but migrating rather than starving in a depleted patch.

After moving foragers fed on the plant in their current cell, if there was one. Foragers were of two types that differed in their feeding behavior. When “restrained” foragers ate, they took only 50% of the plant’s energy. In contrast, “unrestrained” foragers ate 99% of the plant. (We set this parameter at less than 100% so that plants could regenerate after being fed on, rather than being permanently destroyed.)

Unlike the original model (Pepper & Smuts 2000), in this study foragers did not reproduce or die. This simplification made it possible to manipulate population size and the frequency of cooperators as experimental variables.

The schedule

At the start of a run, the program placed plants into evenly spaced square patches with one plant in each cell. Both the size and spacing of patches were variable parameters of the model. The program first placed the specified minimum number of plants into patches, then added plants and empty cells as required to create a uniform square world without partially filled or unevenly spaced patches. Under the standard parameter settings (Table 1), the grid was 112 cells wide and contained $8 \times 8 = 64$ patches, each containing 16 plants (Fig. 1).

After the program created the patches of plants, it placed the specified number of foragers of each type onto randomly chosen cells containing a plant and no other forager. The model was then run for a specified number of time steps. During each time step each agent was activated once in random order, either to grow (if a plant) or to move and eat (if a forager). The standard parameter settings shown in Table 1 were used in each run unless otherwise noted.

Table 1. Standard parameter settings.

	Parameter	Value
Plants	Time steps per run	10,000
	Patch width (grid cells)	4
	Gap width (grid cells)	10
	Minimum number of plants	1000
	Logistic growth rate R	0.2
	Maximum size (energy units)	10
Foragers	Total population size	80
	Population frequency of restrained foragers	50%
	Metabolic rate (energy units per time step)	2
	Feeding restraint (% left uneaten)	1% or 50%




Fig. 1. The distribution pattern of plants (filled squares) on the grid. This figure shows one quarter of the grid created under the standard parameter settings of patch width = 4, gap width = 10, and minimum plants = 1000. To create a uniform square world, the program increased the actual number of plants to 1024 to create 64 patches, each containing 16 plants.

Measuring assortment

To measure assortment one must first define group boundaries. Ideally the groups should be trait groups, defined as the smallest set of individuals whose fitness depends on the traits of other group members (Wilson 1975). However, trait group membership was not imposed as an assumption of this model, and like the real world, the fitness effects of a given act did not have clearly defined boundaries in either time or space. As a proxy for actual trait groups we therefore measured relatedness within patches, which corresponded well to trait groups when patches were small and well separated. Foragers not currently located in a patch were assigned to the last patch they had occupied.

We defined assortment as the deviation of genetic similarity within patches from that expected under random assortment. Genetic similarity was measured by Hamilton's (1972) regression coefficient of relatedness r . Thus the index of assortment was

$$r_a = r - r_s$$

where r_a = genetic similarity due to assortment, r = total (observed) genetic similarity, and r_s = "structural" similarity, or the level of genetic similarity expected under random assortment (Pepper 2000). Because the trait in question, feeding restraint, benefited the actor as well as other members of the group, the appropriate measure of similarity was between the actor and the group it belonged to, including itself. The expected value of this relatedness coefficient under random assortment depends only on population size and average group size, and is defined as

$$r_s = (g-1) / (N-1)$$

where g = number of groups and N = total population size (Pepper 2000). The index of assortment was therefore

$$r_a = r - r_s = \beta(G_w, G_A) - [(g-1) / (N-1)]$$

where β = the regression coefficient, G_A = genotype of actor, and G_w = average genotype of the whole group containing the actor. Genotype was defined as frequency of the restraint allele, which was 0 in unrestrained feeders and 1 in restrained feeders. Positive values of r_a indicate greater similarity within groups than expected, negative values indicate less similarity than expected, and zero is the expected value under random assortment.

Results

Part 1: Can migration patterns generate positive assortment?

To investigate whether patterns of migration between patches could generate positive assortment for feeding restraint, we set all parameters according to Table 1, and performed 10 runs of the model with different random number seeds. The average level of assortment increased rapidly during the runs, from approximately zero (random assortment) to a consistently positive value (Fig. 2). For each run we allowed 1000 time steps for the system to equilibrate, and averaged the assortment level over the last 9000 time steps. This average assortment was positive in every run, and its mean across runs was positive to a high level of statistical significance (mean = 0.045, SD = 0.008, $N = 10$ runs, two-tailed t test, $p < 0.0001$). Thus migration between patches did generate positive assortment.




Fig. 2. Assortment among cooperators as a function of time. The heavy line shows the mean of 10 runs, and light lines show mean \pm one standard deviation. Parameter settings are shown in Table 1.

Part 2: How did migration generate positive assortment?

It is well established that if two types of agents follow different movement rules, this can lead them to segregate, generating positive assortment. We propose a novel mechanism: that even under identical movement rules, segregation can arise through differential effects on the local environment. Specifically, we hypothesized that the positive assortment arose in this model through a mechanism we term “environmental feedback”, which operates when: 1) Individuals differ in their effects on their local environment, and 2) Migration by all individuals is sensitive to local environmental quality.

To determine whether these conditions are indeed sufficient to generate positive assortment, we first had to rule out a potentially confounding factor. In our model, differences in feeding restraint led not only to differing effects on the environment, but also to different thresholds for leaving patches. Both forager types switched from food-seeking to random movement when their food intake fell below the threshold of 2 units, leading to migration out of the current patch. But because restrained foragers ate less of the available food, they reached this threshold when plants were larger (4 units versus 2.02 units for unrestrained foragers).

To eliminate this potentially confounding factor, we repeated the above experiment with one modification. For this experiment only, both restrained and unrestrained foragers moved to the largest available plant if it contained at least 3 units of energy, and otherwise moved randomly. All other rules and parameters were unchanged from the previous experiment, and we performed the same analysis. Again, average assortment was positive in every run, and the mean across runs was significantly positive (mean = 0.041, SD = 0.014, N = 10 runs, two-tailed t test, $p < 0.0001$). Thus positive assortment arose even when all individuals followed identical movement rules.

To further clarify these dynamics, we conducted two control experiments to test whether positive assortment required both of the above conditions hypothesized to be necessary. For the first control we removed differences in environmental effects by setting the level of feeding restraint at an intermediate value of 25% for all foragers. All other parameters remained at their standard values (Table 1). Average values of the assortment index over the last 9000 time steps of each run were statistically indistinguishable from zero (mean = 0.004, SD = 0.008, N = 10 runs, two-tailed t test, $p > 0.05$). Thus positive assortment did not arise without differences between forager types in their effects on their local environment.

For the second control we made foragers insensitive to local food abundance in their movement patterns. To achieve this we modified the movement rule for this experiment only as follows: Foragers surveyed their current cell and the eight adjacent cells, and chose randomly among the unoccupied cells containing any plant. So that foragers would continue to migrate between patches, they moved to this cell with a 90% probability and to a randomly chosen adjacent cell (with or without a plant) with a 10% probability. Again the resulting level of assortment was indistinguishable from zero (mean = 0.014, SD = 0.023, N = 10 runs, two-tailed t test, $p > 0.05$). Thus positive assortment did not arise when foragers were insensitive to their own effects on the local environment.

Part 3: Sensitivity analysis

In this section we investigate the sensitivity of our results to variation in several parameters describing foragers and their environment. In each experiment we varied a single parameter while keeping the others set according to Table 1.

Difference in feeding restraint

The first experiment in our sensitivity analysis extended the result that differences in restraint were necessary to generate assortment. Here we varied the difference in feeding restraint between the two forager types from 0% to 50% while holding the average restraint level constant at 25%. The highest levels of positive assortment resulted from intermediate levels of difference in restraint (Fig. 3).




Fig. 3. Assortment as a function of difference between restrained and unrestrained feeders (mean \pm standard deviation). In all runs the average restraint was 25% (25% of available food left uneaten per time step). Restraint for the two forager types ranged from 25% and 25% (0% difference) to 1% and 49% (48% difference).

Forager population density

In the next experiment we returned to the standard restraint levels of 1% and 50%, but varied the number of foragers from 20 to 400, corresponding to a spatial density ranging from 627.2 to 31.4 grid cells per forager, or a range of 51.2 to 2.6 plants per forager. The highest levels of positive assortment resulted from intermediate population densities (Fig. 4).




Fig. 4. Assortment as a function of population size (mean \pm SD across 10 runs).

Frequency of restrained foragers

We next returned to the standard forager population size of 80, but varied the frequency of restrained versus unrestrained foragers within this population. The level of positive assortment increased monotonically with the frequency of restrained foragers (Fig. 5).




Fig. 5. Assortment as a function of the frequency of restrained feeders in the population (mean +/- SD across 10 runs).

Food growth rate

In this experiment we returned to the standard frequency of 50% restrained foragers, and varied the logistic rate of growth (R) of the plants. The highest levels of assortment resulted from intermediate plant growth rates of about $R = 0.4$. Higher growth rates ($R > 0.5$) produced significantly negative levels of assortment (Fig. 6).




Fig. 6. Assortment as a function of the plant logistic growth rate R . Graph shows mean \pm SD over 10 runs. Asterisks indicate significant differences from zero (t test, $p < 0.05$).

Food distribution pattern

For the final experiment in the sensitivity analysis, we varied the two parameters controlling the distribution of plants: patch width and gap width. Assortment was non-negative under all plant distribution patterns, and highest with intermediate size patches and large gaps (Table 2).

Table 2. Assortment index as a function of patch and gap width. Shaded areas indicate values > 0.02 before rounding.

Patch width	Gap width						
	1	2	3	4	5	7	10
1	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	0.02	0.02	0.02	0.01	0.01	0.00	0.00
3	0.02	0.01	0.01	0.01	0.01	0.01	0.01
4	0.01	0.01	0.02	0.03	0.03	0.04	0.04
5	0.01	0.01	0.02	0.03	0.04	0.05	0.05
6	0.01	0.00	0.00	0.01	0.01	0.01	0.01
7	0.01	0.01	0.01	0.00	0.01	0.00	0.00
8	0.01	0.01	0.01	0.00	0.00	0.00	0.01
9	0.01	0.01	0.00	0.00	0.00	0.00	0.00
10	0.00	0.01	0.00	0.00	0.00	0.01	0.00

Discussion

These results show that positive assortment can arise in an ecologically plausible setting under a fairly wide range of parameter settings. This suggests that Hamilton (1975) was correct in claiming that genetic similarity among group members can arise through mechanisms other than common ancestry, and consequently that inclusive fitness is more general than kin selection. Non-random assortment arose in this model through a process we refer to as “environmental feedback”, which requires both that individual behavior differs in ways that affect the local environment, and that individuals responded to these local differences in their movement patterns.

Exactly how did these conditions generate non-random assortment? Movement patterns were determined solely by the simple rules described above. Under these rules all foragers tended to leave patches with low food abundance and to stay in patches with high food abundance. As a

result, all foragers tended to leave patches containing mostly unrestrained foragers, and to stay in patches containing mostly restrained foragers. Such movements by any forager would tend to reduce its association with unrestrained foragers and increase its association with restrained foragers. When restrained foragers migrated, this put similar types together and generated positive assortment. On the other hand, the same kind of movements by unrestrained foragers tended to put dissimilar types together, generating negative assortment. The balance of these opposing tendencies presumably determined the outcome in any given run of the model. Each parameter of the model had some effect on the sign and magnitude of the resulting assortment pattern.

The sensitivity analyses showed that assortment increased monotonically with two of the six parameters we examined. Large gaps facilitated non-random assortment because they reduced migration among patches. Although some migration was necessary to alter assortment from its initial random state, any excess beyond the necessary minimum moved the system toward a “well stirred” state, in which non-random assortment patterns were hard to sustain. The frequency of restrained foragers was important because, as discussed above, movements by restrained foragers tended to generate positive assortment while movements by unrestrained foragers tended to produce negative assortment. The frequency of movement of each type was determined by their proportion of the total population. In populations dominated by restrained foragers they made most of the moves, and therefore the tendency toward positive assortment dominated the outcome.

For the remaining four parameters, assortment reached positive maxima at particular values. For population density and restraint difference, assortment probably fell off at higher values due to increased rates of population mixing. Population density must be non-zero to produce interactions between individuals, but high population density led to depletion of food everywhere and thus to frequent random movements. Similarly, differences in restraint were necessary to generate differing effects on the local environment, but very low levels of restraint among unrestrained foragers led to rapid depletion of the patches they entered, and thus frequent migration by both forager types. In either case, the high rate of migration by both types removed the non-random assortment that would otherwise have built up.

Assortment was also maximized under intermediate patch size. Non-random assortment required patches above a threshold size because very small patches could hold only one forager at a time, making non-random assortment impossible. As patches became too large, however, assortment returned to near random. According to analyses beyond the scope of this paper, the dynamics of environmental feedback are driven by a small size effect, such that the number of individuals per patch must be small for significant positive assortment to arise.

Finally, assortment was maximized at a moderately low rate of food growth, and actually changed sign to negative assortment at higher growth rates. Why this came about is not clear, and future research will attempt to clarify this phenomenon.

Although we varied only one parameter at a time in the experiments reported here, preliminary investigations have shown that varying more than one parameter can produce non-linear interactions. As a result, our results do not remotely approach an exhaustive search of the parameter space. It is not clear whether other parameter settings would produce more extreme values of non-random assortment, but it seems unlikely that these one-dimensional transects through a six-dimensional parameter space would have chanced upon the most extreme values. And of course, other models could include other parameters that would likely effect the degree of assortment as well. Thus even given our limited attempt at sensitivity analyses, our results should be interpreted mainly in qualitative rather than quantitative terms. The central result is that the mechanism of environmental feedback can produce significantly non-random assortment patterns, including both positive and negative assortment, even when all individuals follow the same movement rules.

Both positive and negative assortment through environmental feedback have important implications for evolutionary theory. As discussed above, positive assortment greatly facilitates the evolution of cooperation, and is essential for the spread of many cooperative traits. Environmental feedback may add to the positive assortment that would otherwise arise through reproduction with limited dispersal of offspring. However, its most important applications may be in understanding the evolution of cooperation between species, where common descent cannot generate positive assortment. For example, many species have two modes of reproduction, local and dispersive. They typically reproduce locally while conditions are good, and produce dispersal forms instead when conditions deteriorate. This mechanism could generate positive assortment for cooperation between species, even when the two species exhibit different forms of cooperation toward each other. It has been shown previously that such correlations in cooperativeness can build up between species, and that they play a role exactly analogous to relatedness through common descent within a species (Frank 1994).

Negative assortment has interesting evolutionary implications as well, not only because it can prevent the evolution of cooperation, but also because it can generate negative relatedness coefficients. Negative relatedness facilitates the evolution of spiteful traits, which spread by reducing the fitness of others even at a cost to the actor (Hamilton 1970). Although the evolution of spite has long been a theoretical possibility, it has received little attention because no plausible mechanism has been known to generate significant negative relatedness. These results suggest a mechanism for generating negative assortment in nature, and the possibility that spite could be more evolutionarily accessible than is usually believed.

How restricted are our basic conclusions to the details of the model we used? In other simulation studies, we have found that positive assortment can arise through environmental feedback in superficially different models. When alarm calling was substituted for feeding restraint as the altruistic trait (Pepper & Smuts 2000), positive assortment still arose. In this case the rate of leaving patches was a function of the local environment not because individuals migrated away when surrounded by “defectors” (non alarm-callers), but because they suffered higher mortality rates. We have also investigated much simpler models showing that whenever the rate of leaving groups varies as a function of the local frequency of altruism, non-random assortment can result (unpub. data). A wide range of altruistic traits could potentially meet such conditions.

Obviously these results do not prove that environmental feedback occurs in nature, or that it generates sufficient positive assortment to drive selection for altruism. Important questions remain about the range of conditions under which environmental feedback will operate. It is also worth noting that the relatedness coefficients resulting from environmental feedback in these simulations were less than 0.1, much lower than the relatedness levels that can arise through common descent (e.g., 0.5 for full sibs or 0.25 for half sibs). However, the results are significant in showing that this novel mechanism can operate, and that the necessary conditions may be fairly general. This suggests that the possibility of positive assortment through environmental feedback in nature deserves serious consideration.

Acknowledgments

We enjoyed helpful discussions of the ideas in this paper with many people, including K. Bassler, D. Goldbaum, T. Kepler, T. Knudsen, M. Lachmann, D. E. Smith, and D. S. Wilson.

Literature Cited

Axelrod, R. & Hamilton, W. D. 1981. The evolution of cooperation. *Science* 211:1390-1396.
Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books, New York.

- Cox, S. J., Sluckin, T. J. & Steele, J. 1999. Group size, memory, and interaction rate in the evolution of cooperation. *Current Anthropology* 40:369-377.
- Eshel, I. & Cavalli-Sforza, L. L. 1982. Assortment of encounters and evolution of cooperativeness. *Proceedings of the National Academy of Sciences USA* 79:1331-1335.
- Frank, S. A. 1994. The genetics of mutualism: The evolution of altruism between species. *Journal of Theoretical Biology* 170:393-400.
- Gilpin, M. E. 1975. *Group Selection in Predator-Prey Communities*. Princeton University Press, Princeton.
- Grafen, A. 1984. Natural selection, kin selection and group selection. Pp. 62-84 in J. R. Krebs and N. B. Davies, eds. *Behavioural ecology: an evolutionary approach*. Sinauer Associates, Sunderland MA.
- Hamilton, W. D. 1964. The genetical evolution of social behaviour. I and II. *Journal of Theoretical Biology* 7:1-52.
- Hamilton, W. D. 1970. Selfish and spiteful behaviour in an evolutionary model. *Nature* 228:1218-1220.
- Hamilton, W. D. 1972. Altruism and related phenomena, mainly in social insects. *Annual Review of Ecology and Systematics* 3:193-232.
- Hamilton, W. D. 1975. Innate social aptitudes of man: an approach from evolutionary genetics. Pp. 133-155 in R. Fox, ed. *Biosocial anthropology*. Malaby Press, London.
- Hart, D. D., Kohler, S. L. & Carlton, R. G.. 1991. Harvesting of benthic algae by territorial grazers: the potential for prudent predation. *Oikos* 60:329-335.
- Maynard Smith, J. 1964. Group selection and kin selection. *Nature* 201:1145-1147.
- Maynard Smith, J. 1976. Group selection. *Quarterly Review of Biology* 51:277-283.
- Minar, N., Burkhart, R., Langton, C., and Askenazi, M. 1996. The swarm simulation system: a toolkit for building multi-agent systems. Santa Fe NM: Santa Fe Institute Working Paper 96-06-042.
- Nowak, M. A. and Sigmund, K. 1998. Evolution of indirect reciprocity by image scoring. *Nature* 393:573-577.
- Pepper, J. W. 2000. Relatedness in group-structured models of social evolution. *Journal of Theoretical Biology* 206:355-368.
- Pepper, J. W. & Smuts, B. B. 2000. The evolution of cooperation in an ecological context: an agent-based model. Pp. 45-76 in T. A. Kohler and G. J. Gummerman, eds. *Dynamics of human and primate societies: agent-based modeling of social and spatial processes*. Oxford University Press, Oxford.
- Price, G. R. 1972. Extension of covariance selection mathematics. *Annals of Human Genetics* 35:485-490.
- Queller, D. C. 1992. Quantitative genetics, inclusive fitness, and group selection. *American Naturalist* 139:540-558.
- Trivers, R. L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology* 46:35-57.
- Wilson, D. S. 1975. A theory of group selection. *Proceedings of the National Academy of Sciences USA* 72:143-146.
- Wilson, D. S. 1977. Structured demes and the evolution of group advantageous traits. *American Naturalist* 111:157-185.
- Wilson, D. S. 1979. Structured demes and trait-group variation. *American Naturalist* 113:606-610.
- Wilson, D. S. 1990. Weak altruism, strong selection. *Oikos* 59:135-140.
- Wilson, D. S. and Dugatkin, L. A. 1997. Group selection and assortative interactions. *American Naturalist* 149:336-351.