

Varieties of Population Structure and the Levels of Selection

Peter Godfrey-Smith

ABSTRACT

Group-structured populations, of the kind prominent in discussions of multilevel selection, are contrasted with ‘neighbor-structured’ populations. I argue that it is a necessary condition on multilevel description of a selection process that there should be a non-arbitrary division of the population into equivalence classes (or an approximation to this situation). The discussion is focused via comparisons between two famous problem cases involving group structure (altruism and heterozygote advantage) and two neighbor-structured cases that resemble them. Conclusions are also drawn about the role of correlated interaction in the evolution of altruism.

1 *Introduction*

2 *Two Kinds of Population Structure*

3 *Objections and Replies*

4 *Particles on a Line*

5 *Conclusion*

Appendix: Neighborhoods and Selection

1 Introduction

Suppose we have an evolving population that is *structured*. Perhaps the population is divided into ‘demes,’ or into shorter-lived social groups. Perhaps we have a population of alleles that are collected into diploid genotypes. Other kinds of structure will be discussed below.¹ Suppose we also know that this population structure *matters* to evolution, in the following sense: the fitness of at least some individuals in the population depends on who they interact with, and the population structure determines who interacts with who.

In both philosophy and biology, considerable attention has been focused on the question of when, in situations like this, the system is best described in terms of multiple ‘levels’ of selection. At least sometimes, it seems natural to

¹ A population can have several kinds of structure at once, but I will generally ignore that possibility for the sake of simplicity.

describe higher-level collections of the lower-level entities as units or targets of selection in their own right.

One aspect of this discussion has been debate about the status of descriptions of evolution in structured populations that do *not* recognize higher levels or units of selection.² Suppose someone seeks to describe an evolutionary process of this kind in terms of selection only at the lower level, and treats the population structure as an aspect of the *environment* in which lower-level selection occurs. Formally, such a description is possible. But is this description legitimate, perhaps even superior, or does it involve some sort of sleight of hand or distortion of the situation? The literature on this question is now large, and has not produced consensus.³ In some of the most recent contributions, the authors have argued strongly for the view that, at least in some crucial cases, the lower-level description *is* deeply misleading. Claims of this kind have been made by Sober and Wilson ([1998]) about behavioral evolution in social groups, by Lloyd ([2005]) about genic selection, and more generally by Sarkar ([2008]).

My aim in this paper is not to wholly resolve this issue. I present an argument that casts light on some famous problem cases, however. My strategy is as follows. I introduce a pair of new cases that are *akin* to some well-known and controversial ones. The controversial cases are the ‘trait group’ model of altruism, and heterozygote advantage in a two-allele genetic model. The new cases are constructed by substituting neighbor-based interactions for group-based interactions in each situation. I argue that in the two neighbor-based cases, a lower-level description of the system is clearly acceptable, and is the *only* description that is acceptable.

The differences between the group-based and neighbor-based cases are then analyzed formally. This comparison is used to draw several conclusions. First, I try to give more precise content to the intuition that *some* kinds of effect of population structure on evolution can properly be described in terms of higher levels of selection, while other kinds should not. The result of this analysis is the formulation of a necessary condition on when a selection process should be described in multilevel terms. This argument draws on work by John Maynard Smith.

A second set of conclusions concerns the right way to summarize what we have learned from recent models about the evolution of altruism. Correlated interaction is widely seen as the key to the evolution of altruism, but this is sometimes presented as a principle *within* multilevel selection theory. This is mistaken, because correlated interaction can be a feature of systems for which multilevel description is not applicable. In fact, a consideration of the

² I will not distinguish between questions about ‘levels’ and questions about ‘units,’ though some authors (e.g. Brandon [1982]) have.

³ Okasha ([2006]) includes a comprehensive survey of work to date.

special features of neighbor-structured populations puts pressure on the familiar distinction between ‘individual’ and ‘group’ selection.

Lastly, I make distinctions among the arguments that might be given against lower-level descriptions of the familiar cases discussed in the literature. Some overly general arguments cannot be sustained, because they would rule out the use of lower-level accounting in the new cases as well as the familiar ones.

2 Two Kinds of Population Structure

In all the model cases discussed in this paper, we imagine an evolving population of ‘particles’ that reproduce asexually. The term ‘particle’ is intended as a neutral description that will cover, in different cases below, a variety of biological entities. Particles come in two types, *A* and *B*. We assume a very large population, nonoverlapping generations, and faithful inheritance of a particle’s type. The four cases treated in the paper differ with respect to the way the population of particles is subdivided or organized.

The first case discussed is a familiar one. The particles are seen as individual organisms that live in social groups that dissolve and reform each generation. This is the ‘trait group’ life cycle that has been extensively investigated by Wilson ([1975], [1980]) and others. In each generation, groups form, selection occurs, and then the groups dissolve at the reproduction step, creating a single ‘pool’ of new individuals from which a new generation of groups is formed. The rule of group formation might either be random or nonrandom.

Suppose that the fitness of both types, *A* and *B*, increases as a function of the number of *A* types in an individual’s group. And further, within any mixed group the *B* type does better than the *A* type. The *A* type is then a kind of ‘altruist.’⁴

In a case like this, it is formally possible to think about fitness in two different ways (Kerr and Godfrey-Smith [2002]). One way is to assign fitnesses only to particles, but do so in a context-sensitive way. A full ‘fitness structure’ for the system will consist in a set of fitness parameters representing the fitness of *A* separately for *each* of the possible group types that an *A* individual might find itself in. The same is done for the *B* type.⁵

The other option is to recognize group-level fitnesses, representing the total outputs or productivities of the various kinds of groups. These group-level outputs are measured in terms of the average or combined productivity of

⁴ This form of altruism has been called ‘individual-centered’ or IC altruism in (Kerr *et al.* [2004]). This form of altruism can be favored under random group formation provided that further conditions are met.

⁵ These are the α_i and β_i parameters of (Kerr and Godfrey-Smith [2002]). Note that these symbols are used slightly differently in the Appendix below.




Figure 1. The trait-group life cycle.

all the particles in the group.⁶ We then also have a second set of parameters representing how the output of a mixed group is divided between the two particle types.⁷

Both representations of the system are possible and each can be translated into the other without loss of information. For now, let us not worry about whether one description is better than the other, in particular cases or in general. Let us also not worry about what this possibility of multiple forms of description implies for causal explanation.⁸ For now, all that matters is that we *can* assign fitnesses to groups, and treat these groups as units in a higher-level selection process. If *A* succeeds over *B*, this might be explained in terms of *A* being favored by a process of group selection, despite its disadvantage within mixed groups (Wilson [1980]; Sober and Wilson [1998]).

A number of writers have explored the question of what features a ‘group’ should have in order to be counted as real and nontrivial (Sterelny [1996]; Sober and Wilson [1998]; Okasha [2006]). Is the mere presence of interaction between

⁶ This is an ‘MLS1’ description, in the sense of Damuth and Heisler ([1988]). Group productivity or fitness is the sum of the reproductive output of the constituent particles. All the multilevel descriptions in this paper are assumed to be of this type.

⁷ These are the π_i and ϕ_i parameters in (Kerr and Godfrey-Smith [2002]).

⁸ In (Kerr and Godfrey-Smith [2002]) we advocate ‘gestalt-switching pluralism.’ *Both* perspectives are valuable, as they ‘package’ information differently. The arguments in the present paper are intended to complement, but not presuppose, our arguments for gestalt-switching pluralism.

members sufficient, or must groups have some level of ‘cohesiveness’? I bracket that issue here; my focus is on some formal characteristics of groups, that affect their role as bearers of group-level fitnesses. These formal characteristics are shared by paradigmatically ‘real’ groups, such as ant colonies, and also the usual examples of ‘marginal cases,’ such as temporary coalitions.

I approach the issue by first looking closely at the role played by groups in the familiar case outlined above. First and most obviously, groups here are *collections* of the lower-level particles. The number of groups in the system is smaller than the number of particles in the lower-level population. The number of groups is also greater than 1. Each particle is a member of only one group.⁹

These features make it possible to think of the groups as forming a higher-level population of competing units that can be assigned fitness parameters. We can treat the productivities of individual particles as contributions to the productivities of higher-level entities. As Wilson has argued, the fact that these groups are ephemeral within the life cycle is not especially important.

In real-life systems, we expect some of these features to be found only approximately. There will be some vagueness of boundaries, and perhaps overlap of groups. I return to the role of approximation later. First I continue the analysis of the idealized cases.

Introducing a more formal language, we can say that groups as found in Case 1 are *equivalence classes* of individuals. If groups are described by means of the relation *x is in the same group as y*, defined on the lower-level particles, we find that this relation is reflexive, symmetric, and transitive. It is an ‘equivalence relation,’ that collects entities into equivalence classes.¹⁰

Once we have collected the particles into equivalence classes, we can describe the system in terms of group-level fitnesses. The ‘Price equation’ can also be used to give a partition of the overall covariance between particle character and fitness into within-group and between-group terms (Price [1972]; Okasha [2004], [2006]). For these formal operations, all we need is *some* means of collecting the particles into equivalence classes. However, it is widely agreed that this should not be done in an arbitrary way. Groupings should be recognized on the basis of biologically real relationships between particles. That is to say,

⁹ As D. S. Wilson has emphasized, a single particle might be a member of several groups that are recognized in the analysis of different selection processes. A person can be a member of a family, a church, and a trade union, all at the same time, and these groups are not hierarchically organized. But if we are analyzing competition between *one* of these sets of groups (churches, for example) then it should be approximately true that each individual is a member of only one of these collective entities—at least that there be many fewer churches than individuals. So my claims here concern the analysis of a particular selection process, in which some particular form of population structure is relevant. See also footnote 1.

¹⁰ The partitioning of a population into groups might occur several times over a generation, so an individual’s group members might vary over time. Then the consequences of all these episodes of group formation must be taken into account in a model. But any particular group in which an individual finds itself, however temporarily, is an equivalence class.

the equivalence relation—*x is in the same group as y*—is not supposed to be primitive and unanalyzed; it is supposed to have a basis in some other relation, or set of relations. If we can express the criterion determining groups in terms of some single underlying relation between particles, then this must be an equivalence relation, given the structure that it is supposed to explain.

It is common in the modeling literature to make the grouping according to which individuals have an effect on other individuals' (absolute) fitnesses (Wilson [1980]; Uyenoyama and Feldman [1980]; Michod [1982]; Wade [1985]; Sober and Wilson [1998]).¹¹ This sort of criterion works well in situations like Case 1, where the fitness structure implies that each individual benefits from the presence of all the altruistic individuals within its group but from no individual outside it.¹² This criterion for determining groups does not involve the use of a single equivalence relation (because the relation is not reflexive), but it can be modified so that it does. Essentially, criteria of this kind make a grouping of individuals in terms of the relation, *x has its fitness affected by the character of y*. In scenarios like my Case 1, this relation is reflexive, symmetric, and transitive. So in cases like this we can simply look at this pattern of relations between particles, and use these relations to partition the population into a set of nonarbitrary groups around which lines can be drawn as shown in Figure 1.

Models are idealized, of course, and in real life a biological group need not meet these requirements exactly. There might be some overlap between groups (which is not possible with equivalence classes). But the relevant concept that is being *approximated* in those cases is the concept of an equivalence class. Here it is significant that the most popular mathematical framework for representing multilevel selection, the Price equation, requires that this condition hold exactly.

Although the modeling tradition has often assumed a grouping of individuals on the basis of effects of each individual on the fitness of all others, this criterion can be problematic when applied to empirical cases. For example, it does not apply to the celebrated cases of group selection involving eusocial insects. (This point was made by an anonymous referee.) Here, the character of a sterile worker affects the fitness of the queen but not vice versa, and the worker's character has no effect on its own (direct) fitness. So *x has its fitness affected*

¹¹ Wade and Michod both endorse this definition from Uyenoyama and Feldman. 'A group is the smallest collection of individuals within a population such that the genotypic fitness calculated within each group is not a (frequency dependent) function of the composition of any other group' (Uyenoyama and Feldman [1980], p. 395). More informally, Sober and Wilson ([1998]) say that a group is a collection of individuals that all influence each others' fitness with respect to a certain trait, but do not influence the fitness of those outside the group (p. 92). As the text notes, it is important that fitness here be understood as *absolute* fitness, or the test becomes trivialized.

¹² I follow common practice in using a simple counterfactual test for the assessment of whether one individual 'affects' another's fitness, even though Rinard ([unpublished]) has shown that this leads to trouble in some cases.

by the character of y is not an equivalence relation that can be used directly to form equivalence classes of individuals, even if we ignore such things as aggressive encounters across colonies. An ant colony is still an equivalence class of individuals (allowing some idealization), however, and what this case shows is that some other relation defined on individuals must be used to establish the grouping. Common ancestry and physical interaction are obviously appropriate criteria here, but the modeler's abstract notion of fitness-affecting interactions might also be applied, albeit in more complex ways. If directed edges (lines with arrows) are drawn between individuals x and y whenever x 's character affects y 's fitness (in a sufficiently systematic manner), and then the arrows are erased and indirect connections are also stipulated to be sufficient for connectedness, we might be able to see an ant colony as a bounded network of direct and indirect systematic effects of individual character on fitness. The relation x is connected to y is then an equivalence relation that partitions the population into separate colonies.

Aside from this complex case, it can be objected that in natural populations, effects of one individual on another's fitness are far more widely distributed than the modelers' criterion for group membership seems to suppose. Even if we follow Wilson ([1980]) and relativize group membership to a trait of interest, subtle effects of one individual on another's fitness may extend across what may seem to be obvious group boundaries. The practice of modelers may best be understood by treating them as imposing an idealized binary distinction on a more graded reality.¹³ In the remainder of this paper, I assume the availability of a simple distinction favored by modelers, between individuals that do and don't affect each others' absolute fitness. But I emphasize that this involves idealization and also is not the only relation defined on individuals that can be used to partition a population into equivalence classes.

I now introduce a second case. In Case 2 we find the same kind of lower-level population of particles, but a different form of population structure. The population is now arrayed on a two-dimensional lattice. As in the first case, the life cycle includes the periodic formation and dissolution of the population structure. Each generation begins with particles in a common 'pool,' from which they settle in a two-dimensional array (perhaps randomly, perhaps not). The reproduction step in the life cycle dissolves the array and generates a new pool.

Selection occurs on the lattice, and the fitness of a particle is affected by the characteristics of its neighbors—specifically, by its immediate four-member N-S-E-W neighborhood or 'Von Neumann neighborhood.' Models of this kind

¹³ I am indebted to an anonymous referee and to Rasmus Winther for emphasizing this point to me.




Figure 2. A model with ephemeral two-dimensional spatial structure.

can also use the eight-member ‘Moore neighborhood’ that includes diagonal neighbors, and there are other neighborhood concepts as well (Skyrms [2004]).

To make the case as similar to the first one as possible, suppose again that each type does better in proportion to its number of *A*-type neighbors, but *B* does better than *A* in any given neighborhood. So again, *A* is a kind of ‘altruist.’ (See the Appendix for further details.)

Case 1 with groups and Case 2 with neighborhoods are clearly quite similar. And although they do not have as long a history in the literature, models with two-dimensional spatial structure and interactions between neighbors have become prominent in recent years—the models in (Alexander and Skyrms [1999]) and (Skyrms [2004]) are elegant examples. These models do not usually use an ‘ephemeral’ spatial lattice that is dissolved and reformed each generation. Usually, the spatial distribution persists and evolves across generations. I introduce an ‘ephemeral’ lattice to make Case 1 and Case 2 as similar as possible; the consequences of this will be discussed below.¹⁴

It is initially apparent that there are no group divisions in this population when it is located on the lattice. That suggests that groups cannot function as units of selection in such a case, even though this is a system in which population structure matters to evolution. I will argue that this is indeed the right conclusion to draw, but it will take some time to make the case, as some other authors have sought to defend different views.

¹⁴ This life cycle is not merely of theoretical interest. It might apply quite well to some marine invertebrates, like some cnidarians, that are dispersed in reproduction but then settle on a two-dimensional surface and interact with neighbors.

As in Case 1, this is a case where the fitness of each particle is affected not just by its own character, but also by the characters of a subset of the other particles in the population. In Case 1, this was used as the criterion for recognizing groups in the population, via the fact that the relation *x has its fitness affected by the character of y* was an equivalence relation on the population of particles. The result is a partition into nonarbitrary groups that can be assigned group-level fitnesses and/or analyzed with the covariance techniques of the Price equation.

In Case 2, some parts of the same story can be told, but parts of it cannot. Once again, we can describe the population in terms of a relation between particles, *x has its fitness affected by the character of y*. But that is not an equivalence relation in this case, because it involves neighborhood relations that are not transitive. If the character of individual *x* affects the fitness of *y*, whose character affects the fitness of *z*, this does not imply that *x* affects the fitness of *z*. Consequently, this relation cannot be used, at least directly, to partition the population into a number of equivalence classes. Each individual determines a unique neighborhood; we have as many neighborhoods as individuals. The only way to partition the population into a range of groups, in a way that yields fewer groups than particles, is to do so arbitrarily.¹⁵ These groups could be of any size and shape.

The most straightforward ways of describing a process of competition between groups, which were employed in Case 1, cannot be applied here. Disregarding the option of arbitrary divisions, we cannot treat each particle's reproductive output as a contribution to the output of a single collective entity that competes with others, in a higher-level selection process. Following the definitions of groups used by the modelers cited earlier, we find there is no division of the population into groups such that everyone inside each group affects the fitness of everyone else inside, and no one inside any group has their fitness affected by the character of an individual outside the group. The partitioning of group-level and individual-level covariances between character and fitness seen in the Price equation is not possible here. (Alternatives to Price will be discussed below.)

It might be thought possible to treat each 'focal particle plus its neighbors,' a cross-shaped object, as a group that competes with other groups. Let us think about what that would involve. There is the same number of these 'groups' in the system as there are particles. Any particle is a member of five different groups of this kind (one for its own neighborhood, and one for each of its neighbors, now taken as focal particles). All the nonfocal particles in any given

¹⁵ Sterelny, in response to Maynard Smith, briefly mentions this possibility in his work ([1996]), without saying much for or against it. I don't know of others who have considered or defended the possibility explicitly.

group have their fitness affected by the character of particles outside the group. Indeed, a nonfocal particle within a given group has its fitness affected by the focal member of the group and by three particles outside that group. So for *most* of the groups that an individual is counted in, *most* of the interactions that affect its fitness are with individuals outside the group.

So suppose we sought to represent evolution in such a case in terms of between-group and within-group processes. In working out the between-group process, each particle's reproductive output would be quintuple-counted; it would contribute to the fitness of five different groups. Each particle's output would also have five different roles in the within-group process as well. If we ask whether some particle is doing well or badly in within-group competition, this assessment of individual success is now made five different times against five different (but overlapping) sets of competitors. Again, most of these comparisons are with particles that do not directly affect the fitness of the one we are concerned with. It may be possible to use this convoluted method to calculate the frequency for the *A* type in the next generation, but such an analysis seems entirely at odds with the structure of the case being analyzed.¹⁶

Population structure *matters* in this second model; the distribution of neighborhoods encountered by individuals affects what will happen. Random and nonrandom rules for the allocation of particles to the lattice will have different effects, for example (see the Appendix). But this population structure is not of a kind that yields a set of nonarbitrary higher-level units that can reasonably be said to *compete*, and that can be plausible bearers of fitnesses. To use an older terminology, neighborhoods in this case cannot be seen as anything like a set of higher-level 'interactors' (Hull [1980]; Lloyd [2001]) that interact with their environment as wholes and cause the differential reproduction of lower-level particles. Instead, this purely neighbor-based form of population structure is a kind that is properly treated as context, as environment.¹⁷

This line of argument is not original to me. It is due, in its core, to John Maynard Smith ([1964], [1976], [1987], [2002]). Maynard Smith tended to make the argument informally and telegraphically, however, and did not introduce the concept of equivalence classes as the crucial one. Maynard Smith often presented the argument as a point about kin selection, and its relation to group selection. He notes that if we have many trees of a single species in a forest, in a two-dimensional array, neighbor-interactions may affect fitness, and do so in

¹⁶ As Ben Kerr pointed out to me, it is worth noting here that a simple 'connectedness' criterion for group membership, introduced earlier to deal with social insects, will not help the attempt to give a multilevel description of this case, because it would turn the whole population into a *single* group, making group-level differences in fitness impossible.

¹⁷ This is not to say that *no* models with neighbor structure can have groups of the right kind. For example, models with neighbor-interactions and local reproduction into regions of empty space *could* give rise to fairly discrete nonarbitrary groups. The model in (Mitteldorf and Wilson [2000]) the potential to give rise to situations of that kind.

a way modified by relatedness between nearby trees, even though there are no discrete groups in the population. So not all cases of kin selection should be seen as group selection. I have sharpened the argument up with the appeal to equivalence classes in contrast with other relational structures, and I treat the argument as having a more general application than the proper interpretation of kin selection (as did Maynard Smith in [2002]).

Maynard Smith's argument seems to have had limited impact. (Some reasons for this will be discussed shortly.) But the phenomena he describes have certainly been noted by others. In his original paper on group selection, Wilson ([1975]) says that 'trait groups' in his sense are sometimes 'discrete' and sometimes 'continuous.' In the latter case, each individual is the center of its own trait-group (p. 143). Wilson's strategy was to give an explicit model for the discrete case, and claim that the continuous case is essentially similar. In his subsequent work, Wilson notes that the continuous case (as he presents it) is hard to model analytically, and he uses computer simulations to argue that the continuous case yields similar behavior to the discrete one, and supports similar conclusions (Wilson [1977], [1980]; Middeldorf and Wilson [2000]).

So it is in the interpretation of the relationship between the 'discrete' and 'continuous' cases that I differ with Wilson. Wilson has not attempted to give an actual model of the continuous case that utilizes a concept of group-level fitness. Because of the additional mathematical complexity of the continuous cases he discusses, Wilson studies them with simulations, and imports a verbal *description* of the results that comes from the *discrete* case. But the description of the continuous case in terms derived from the discrete one is, as I argued above, very problematic.

This handling of the relations between the more tractable and intractable cases is characteristic of the field. In fact, I suggest that people's thinking about these theoretical issues has been affected in interesting ways by methodological factors. There is a great deal of work being done on spatially explicit models at present.¹⁸ These spatial models do not usually generate analytical results that apply over more than one generation; most of this work is simulation-based. This is because in most of these models, the distribution of neighborhoods encountered by each type *co-evolves* with population itself, and does so in complex ways. These are treated as 'viscous' populations, to use Hamilton's term ([1964], [1975]). In these models there is usually no rule, akin to a Hardy-Weinberg rule, that will generate a new population structure in a regular way from the new generation of particles resulting from the previous

¹⁸ In addition to the works cited above, see (Nowak and May [1992]; Wilson *et al.* [1992]; Alexander [2003]; Werfel and Bar-Yam [2004]; and Leiberman *et al.* [2005]). This work also connects to an independent tradition of work on 'cellular automata,' as noted by Nowak and May.

round of selection. An initially random distribution does not stay random for long.

Most *analytical* work on structured populations assumes groups that are represented as equivalence classes. And, I suggest, the theoretical principles that people find compelling tend to be principles *derived from the analytical work*. If one's theoretical picture is derived from work that assumes equivalence classes, then it can be natural and tempting to say that when population structure is present and affects fitness, we must have more than one level of selection. My argument is that to the extent that equivalence classes cannot be recognized (at least approximately) in the population structure, this conclusion does not follow.

Once we see the role of these methodological factors, it is interesting to reconsider some famous earlier discussions. A key example is (Hamilton [1975]). Hamilton's paper is now often read as a dramatic endorsement of multilevel selection (Sober and Wilson [1998]). Parts of the paper certainly have this character, especially his use of the Price equation. However, in a brief but crucial passage Hamilton notes that his own preferred approach can handle cases in which altruism can be favored without groups being present. The 'inclusive fitness' approach, he says, 'can deal with an ungrouped viscous population where, owing to restricted migration, an individual's normal neighbors and interactants tend to be his genetical kindred.' He then notes that even kinship itself is not essential to the inclusive fitness approach as he conceives it. '[K]inship should be considered just one way of getting positive regression of genotype in the recipient [of altruistic behaviors], and . . . it is this positive regression that is vitally necessary for altruism' (all p. 337). This passage appears at the end of a modeling treatment in which equivalence classes are assumed, and group effects are analyzed via Price. This suggests that Hamilton may have treated the multi-level framework as a convenient and elegant framework for representing one particular manifestation of the phenomenon he was interested in, but not as representing the whole or the essence of it.¹⁹

The main conclusion to be drawn from this section can be expressed as follows. It is a necessary (not sufficient) condition on the existence of higher levels of selection that there be a nonarbitrary partition of the lower-level

¹⁹ Perhaps some of Wilson's earlier discussions ([1977], [1980]) can be read the same way (e.g., [1977], p. 183). It may seem there that Wilson has his eye on a general phenomenon that in *some* cases can be compactly represented in terms of group selection, but not in all cases. However, Wilson's later discussions are more emphatic and unequivocal about the essential role of selection on groups, and it is these that I take issue with. Here is an example: 'For many years, genealogical relatedness was the guiding light for biologists interested in the evolution of altruism. It has become apparent more recently that *variation among groups is the essential ingredient*, which can be accomplished by more than one mechanism.' (emphasis added; Wilson and Dugatkin [1997], pp. 348–9). See also (Sober and Wilson [1998], p. 57): 'the theories that have been celebrated as alternatives to group selection are nothing of the sort. They are different ways of viewing evolution in multigroup populations.'

population into equivalence classes, or at least an approximation to this condition. Following the lead of the modelers discussed, I have focused on criteria for recognizing groups that involve direct effects of one individual on the fitness of another. But there are other criteria that would meet the formal criterion defended here. There will also be ways of using several biological relations in tandem to achieve a grouping of the right kind.

3 Objections and Replies

In this section I reply to two possible objections. Both try to give theoretical reasons for the counterintuitive conclusion that group selection is operating in the neighbor-structured model, Case 2.

First, in the preceding paragraphs I used the requirements associated with Price's covariance-based method for detecting group selection to support the claim that neighbor-structured cases cannot be handled in multilevel terms. Many other tools used in analyzing group selection require equivalence classes as well.²⁰ But one important alternative to the Price approach is 'contextual analysis' (Damuth and Heisler [1988]; Okasha [2004], [2005]). Contextual analysis is based on multiple regression techniques; it assesses group selection by measuring the partial regression of particle fitness on 'group character,' controlling for particle character. The original applications of contextual analysis to evolutionary questions assume the presence of equivalence classes, but contextual analysis can easily be modified to deal with cases of neighborhood structure, as opposed to group structure, as Okasha ([2005]) and Kerr ([in correspondence]) have noted. This can be done by substituting 'neighborhood character' for 'group character'; the presence of overlaps between neighborhoods would not be a problem, as it is for Price. So an available—perhaps superior—way of representing group selection does not require equivalence classes.

Contextual analysis is clearly a useful tool. The harder question is what exactly it measures. When taken as a test for group selection, contextual analysis is taken by many to have a serious flaw. It recognizes a positive group-level term in cases known as 'soft selection,' where there is frequency-dependent selection within groups, but the groups themselves all have exactly the same productivity (or, in some models, fixed productivities that may differ among themselves but are not affected by the trait distributions within the groups).²¹ This shows that

²⁰ As Okasha (personal communication) notes, this is true of a one-way ANOVA framework that was often used before Price's equation became popular. It is also true of the usual framework used to analyze 'interdemic' forms of multilevel selection.

²¹ The observation was originally made by Goodnight *et al.* ([1992]), who do *not* regard it as a decisive problem for contextual analysis. Rice ([2004]) regards it as showing that contextual analysis does not test for higher-level selection in the usual sense (p. 325). Okasha ([2006]) regards it as problematic but as something to be weighed against other advantages.

contextual analysis is, at bottom, a test for detecting whether or not a given kind of population structure *matters* to evolution. It is not a test for whether or not some groups are fitter than others, or whether a certain trait is giving some groups an advantage over others.

For some authors, including Okasha, contextual analysis remains a superior approach to the detection of group selection despite this shortcoming.²² If one is convinced that the core of the group selection problem consists in distinguishing real from spurious *effects* of group structure, then it will be natural to use contextual analysis rather than Price. This is a minority position, however, as most theoreticians want to hold fast to the idea that group selection does require differences in the reproductive output of groups (Wade [1985]; Sober and Wilson [1998]; Rice [2004]). That point has usually been made within a modeling framework in which a grouping of individuals into equivalence classes was assumed. The question has not been assessed in a context in which equivalence classes are absent, even in approximation, like my Case 2. An explicit consideration of these purely neighbor-structured populations adds further support to the majority position. This is because when we note the applicability of contextual analysis in a system with no group divisions at all, it is clear how distant the term ‘group selection’ is from its antecedent meaning, when one treats contextual analysis as testing for it. A person could say that they explicitly intend to understand the phrase ‘group selection is acting’ to mean the same as ‘population structure of some kind is making a difference to evolutionary outcomes,’ but this is now very far indeed from the traditional concepts of group-level advantage and group fitness.

The second objection I will discuss is closely related—largely a variant on the same line of reasoning. In one of the most detailed and useful treatments of neighbor-structured populations, Len Nunney ([1985]) suggests that we identify group selection with ‘a process by which a trait spreads or is maintained in a population because of the differential reproduction of genotypes that arises from the positive association of individuals exhibiting the trait’ (p. 221). In addition, Okasha ([2005]) uses Nunney’s analysis as the basis for a qualified defense of a view of group selection based on a modified form of contextual analysis, under which there is group selection if and only if there is a partial

The term ‘soft selection,’ as applied to this sort of population-structured evolution, originates with Wallace ([1968]).

²² The main argument for regarding contextual analysis as superior is the fact that it does not have a ‘false positive’ problem that the Price approach has. The Price approach recognizes group selection in some cases where group-level fitness differences are apparently mere byproducts of lower-level fitness differences. Concern with these byproduct problems goes back to Williams ([1966]) and has been discussed in detail by Sober ([1984]) and Nunney ([1985]). See Okasha ([2006]) for extensive discussion.

regression of particle fitness on ‘neighborhood character.’²³ My main reply to the contextual analysis approach is above, but I will make some extra comments about Nunney’s own definition.

Nunney’s proposed definition of ‘group selection’ is perhaps the most unusual one ever suggested. It does not even mention the role of groups, but attributes any outcome to ‘group selection’ if it arises via positive correlation between types. An immediate contrast can be drawn with the passage from (Hamilton [1975]) quoted earlier, according to which correlation is indeed a key theoretical concept, but correlated interaction is seen as possible within both group-structured populations and ‘ungrouped viscous’ populations. The underlying theoretical picture is similar, but the terms used to describe it are different.²⁴

Why does Nunney opt for his new and unusual meaning of ‘group selection?’ One reason is that he thinks there are good arguments against more familiar analyses. But another reason is that Nunney is guiding his usage with the idea that the contrast between ‘individual selection’ and ‘group selection’ should be sharp, and the two kinds of selection should generate different evolutionary outcomes. In particular, group selection is what maintains traits that are not advantageous to individuals. Via his own work and others’, Nunney accepts the conclusion that groups in an ordinary, pretheoretic sense need not be present for altruism to be favored, provided there is correlated interaction. But he wants to hold onto the idea that altruism is the product of something other than individual selection. The result is a definition in which it is made impossible in principle for group selection to operate if group formation is random, no matter how discrete, cohesive, and ‘real’ the groups are, and no matter how large the differences in group-level output.

It is difficult, and often unwise, to argue against an explicit stipulative definition that is offered in full awareness of its unorthodoxy. So one possible conclusion to draw here is that we should grant Nunney his concept, give it a new label, and simply note where it converges and diverges from more orthodox concepts. However, in this case I think a positive argument can be given against some of the motivations that Nunney offers for his treatment of the term. Nunney is guided, as I said, by a desire to retain a sharp contrast between the evolutionary roles of individual and group selection. I suggest that this is a very weak argument. The apportioning of controversial phenomena such as altruism between ‘group selection’ and ‘individual selection’ should

²³ As it happens, Okasha ([2005]) defines the neighbor approach in a way that assumes discrete groups are present. ‘An individual’s neighbours are defined as all the organisms in its group *except itself*.’ But this could easily be modified.

²⁴ For an especially useful discussion of the role of correlation per se, with other connections to the philosophical literature, see (Skyrms [1994]). A view of altruism emphasizing correlation per se rather than group structure was also sketched, and later abandoned, by (Sober [1992]). See (Frank [1998]) for a detailed and relevant treatment with an emphasis on kin selection.

be the *upshot* of a satisfactory analysis, not a prior constraint that imposes a set of otherwise problematic definitions at the outset. In advance of a detailed analysis, why are we so sure that a particular phenomenon should be allocated to one explanatory factor or the other? And why are we so sure that some processes could not qualify as both ‘individual’ and ‘group’ selection at once?

Further, as Nunney’s own work shows, the concept of ‘individual advantage’ is itself less clear than is often supposed. Sometimes questions of individual advantage are best addressed via a ‘mutation test’ that holds an individual’s background context fixed and imagines a change in the individual’s character, rather than comparing two individuals of different character within the same group. But in other cases, such as that of an unstructured population, the simple comparison of individuals in their common actual setting is appropriate. Kerr and Godfrey-Smith ([2002]) have also argued that the concept of ‘individual fitness’ has more variants than is usually supposed, even leaving aside the controversial case of ‘marginal’ fitnesses. The lesson from this and other work is that the concept of ‘individual selection,’ and its relatives, are somewhat loose and ambiguous ones. Consequently, the motivation for Nunney’s unusual definition of group selection is not at all compelling.

A more general conclusion can be drawn from this consideration of Nunney’s framework. People sometimes object to the individualist description of neighbor-structured populations because it has become clear that strong forms of altruism can evolve in such populations. Surely altruism involves individual-level disadvantage? Some treatments of altruism even make this true by definition. So how can a system with only ‘individual selection’ sustain altruism? My response to this line of thought follows the reply to Nunney described in the preceding text. First, definitions of a phenomenon that constrain its possible explanations by *fiat* should be avoided if possible, and here it is easy to avoid them (Kerr *et al.* [2004]). Second, it might be true that some narrow existing sense of ‘individual selection’ makes it impossible for individual selection to sustain altruism. But the new spatially explicit models that have been cited in this section are expanding our conception of what ‘individual selection’ is. Perhaps it would be more accurate to say that these new models are making some of the traditional contrasts in this area obsolete. We see this, in a way, in the growing realization that correlated interaction is a crucial unifying concept in this area, and one that has both individualist and nonindividualist aspects. We can expect a better framework to be developed as the models evolve. In the meantime, it is a mistake to argue that because we ‘know’ that altruism is at odds with individual selection, any model in which altruism evolves must include group selection.

The next section connects these considerations to the literature on genic selection.

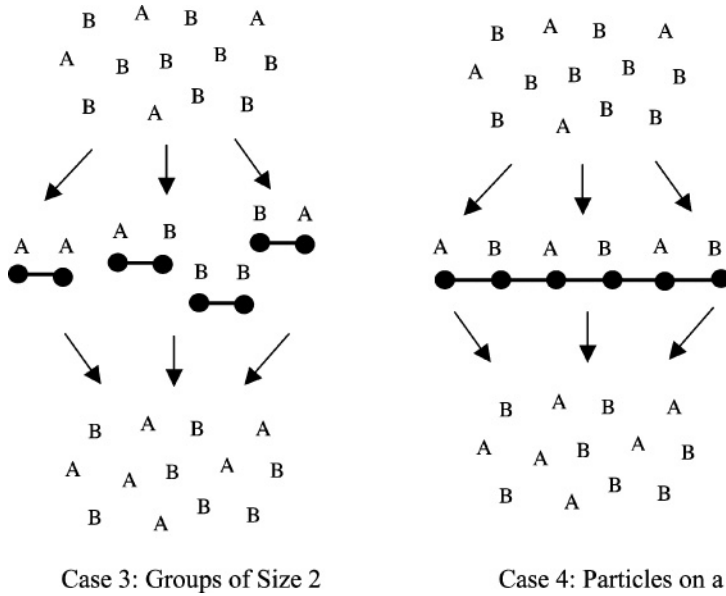


Figure 3. Two additional cases.

4 Particles on a Line

The model in ‘Case 2’ uses neighborhoods in two spatial dimensions. I now introduce another contrasting pair of cases, one with equivalence classes and one with neighbors in one spatial dimension.

Case 3 is like Case 1, except that particles form groups of two rather than five. In Case 4, the particles occupy places on a one-dimensional environment with a well-defined neighborhood relation but no group boundaries.

We suppose that each particle’s fitness in Case 4 is affected by its own character and by the character of the neighbor to its *right*, but not by that of the neighbor to its left. We have context-sensitive fitnesses, as before. And once again, the kind of population structure that affects fitness does not yield a division of the population into groups (equivalence classes) based solely on those interactions. There are as many neighborhoods as there are particles in the population. In Case 2 in the previous section, this was achieved by making the relation between particles that affects their fitness *symmetric but nontransitive*. That method is not possible in the case where each particle’s fitness is affected by only one other particle, so in Case 4 the same result is achieved by making the fitness-affecting relation *asymmetric*.

If we suppose there is a dissolving of the line at the end of each generation, and some fixed rule for the assortment of the new generation of particles into this one-dimensional environment, the resulting system is easy to

describe mathematically (see the Appendix). Each type of particle is described with two fitness parameters, representing its fitness in each of its two possible neighborhoods. The settling of particles on the line may be either random or nonrandom. If it is nonrandom, a formula for change must track not only the overall frequencies of the types, but a set of *conditional frequencies of neighborhoods*. These describe the frequency of *A*-type neighborhoods conditional on the state of the focal particle whose neighborhood this is. These conditional frequencies were discussed as an optional accounting tool for diploid genetic models in (Godfrey-Smith and Lewontin [1993]). Here, they are more than optional; there is no way to divide the line into pairs (equivalence classes) without imposing arbitrary boundaries. So these conditional frequencies are by far the most natural way to describe the role of this kind of population structure.

Let us return to Case 3, which has a population divided into groups of size two, for a moment. As is well known, a trait-group model with groups of size two can be interpreted as a population genetic model of a diploid population, at one locus.²⁵ If the mixed pairs *AB* have the highest fitness of the three kinds of pairs, then it is a model of heterozygote advantage. The fitness of the *A* type is context-dependent; it is higher in the company of *B* than it is with its own type. It is possible to describe this case in terms that assign fitnesses only to particles, using the same kind of accounting that was used for Case 4, the case with particles on a line. But Case 3 can also, of course, be described in terms of fitnesses that are assigned to pairs as wholes, and that is the usual practice. Further, the case of heterozygote advantage is often seen as one demanding analysis in terms of selection at the level of diploid genotype or whole organism, not selection on the level of the lower-level alleles (Sober and Lewontin [1982]; Lloyd [2005]). That might be the right conclusion in this particular case, but what is the argument behind it?

The comparison between Cases 3 and 4 heads off one possible line of argument. It cannot be that the lower-level description of Case 3 is bad because of some general flaw in the form of description that uses ‘contextual fitness of particles,’ or context-sensitive fitnesses at the lower level. This can be seen from the fact that this form of description is the *only* reasonable approach in Case 4, which is in many ways similar to Case 3. Further, it cannot be argued that this kind of context-dependence of fitnesses always implies the existence of a higher level of selection. Context-dependent fitnesses can exist when candidates for higher-level selection do not exist in the system at all.

So some possible ways of objecting to a lower-level description of Case 3, the case with groups of size two, are rejected. This does not show that *in the specific case of genes and genotypes*, there are no problems with a gene’s eye

²⁵ Here, a population that allows selfing.

view. There may be special features associated with the details of that case, and its particular causal structure, that tell against the low-level description.

With the aid of these points, we can see an important division within the lines of argument developed by Sober and Wilson, Lloyd, and Sarkar against various lower-level descriptions of processes that admit of multilevel representation. Sometimes these critics seem to be arguing for a highly general claim, of something like the following form: 'The kind of "contextual" information used in lower-level descriptions of evolution in structured populations *just is* information about higher-level units or levels of selection, differently expressed.' That is the claim that my arguments here are intended to oppose.²⁶ But on other occasions, the arguments made by these critics are based on the causal details of particular crucial cases. Those claims, which must be assessed separately, are not affected by the arguments given here.

Consider the case of (Lloyd [2005]). Lloyd's main focus is the specific case of genes and genotypes. Here, she thinks that a lower-level accounting using context-sensitive allelic fitnesses amounts to a smuggling-in of higher-level information. In the case of genes and genotypes, genes *are* collected together into equivalence classes, at least in the context of the models she is discussing, and their causal roles are tightly integrated in a particular way. Insofar as Lloyd's arguments are based on the contingencies of the case, they are not affected by the points made here.²⁷ The claim to be rejected, again, is the claim that any systematic effect of context on fitness implies the presence of a higher level unit of selection.

5 Conclusion

My primary aim in this paper has been to discuss the roles of some different kinds of population structure for debates about multi-level selection. Most models of structured populations that have inspired the formulation of general theoretical principles assume a partition of the population into equivalence classes. In this background, it has perhaps been natural to suspect that there is some close link between population structure *affecting fitness* and selection operating at a higher *level*. The case of population structures that do not generate equivalence classes of lower-level particles shows that this cannot be true in general. So any argument against lower-level descriptions of particular cases (a gene's eye view of heterozygote advantage, individualism about game theory, etc.) must be based on something other than this general principle.

²⁶ Perhaps Sarkar ([2008]) is the most explicit, in making a general claim of this kind.

²⁷ I do not think Lloyd's 2005 arguments suffice to show the inviability of a properly formulated pluralist view, but that is a separate issue. For discussion of pluralist options, see (Waters [2005]) and (Kerr and Godfrey-Smith [2002]).

Second, I have used a formal distinction between two kinds of relational structures to give a necessary (not sufficient) condition on when population structure can support a description in terms of higher levels of selection. My proposal is that it is a necessary condition on the existence of higher levels of selection that fitness-affecting interactions, or other biological features of the system, partition the lower-level particles into equivalence classes. An approximate meeting of this condition might often be enough, but this requirement is far from trivial, as there are certainly cases (like my Cases 2 and 4) that do not approximate it at all.

The general picture that follows from this proposal is as follows: the evolution of a population is often dependent on the population's embedding in a 'relational structure' of some kind. That feature is seen in all the four cases discussed here. But these structures come in many varieties. The language of multilevel selection theory, when used literally, only applies when the relational structure in question generates equivalence classes, or reasonable approximations to them. Other theoretical distinctions of this kind might be made as well, by attending further to different kinds of relational structure.²⁸

Lastly, I will summarize conclusions that can be drawn regarding the evolution of altruism. One of the most important results from recent work on neighbor-structured populations has been the demonstration of how various kinds of cooperation and altruism can evolve in these models. There has sometimes been a temptation to assimilate these phenomena, at least informally, to cases where there are groups present. Once we understand the special features of neighbor-structured populations, and see them as important in their own right, this assimilation becomes unnecessary and misleading. The focus of work can then turn to the question of what *difference* it makes to evolution when a system contains genuine bounded groups, as opposed to other kinds of population structure, and the question of how such collective entities evolve. This is one aspect of the problem of explaining the 'major transitions' in evolution.²⁹

Appendix: Neighborhoods and Selection

This appendix gives a more formal description of some features of evolution in neighbor-structured populations, with particular attention to altruism. It draws on (Nunney [1985]) and unpublished work by Ben Kerr.

Fitness in neighbor-structured populations can be described with modified versions of the α and β parameters used for group-structured populations in

²⁸ Alexander's recent work looks at some of these additional kinds of structure (e.g. [2003]). See also (Leiberman *et al.* [2005]).

²⁹ See (Buss [1987]; Maynard Smith and Szathmary [1995]; and Michod [1999]).

(Kerr and Godfrey-Smith [2002]). Let α_i be the (absolute) fitness of an A -type individual with i neighbors of the A type, and β_i be the fitness of a B -type individual with i neighbors of the A type. Let p be the frequency of the A type in the population and n the number of neighbors whose character affects the fitness of any individual. The population is assumed to be infinitely large, asexual, and with nonoverlapping generations.

Some standard definitions of altruism cannot be applied to a neighbor-structured case, as they require a comparison of the fitnesses of two individuals within a common group, and/or a comparison of group-level fitnesses in which individuals outside a given group do not affect the fitnesses of those inside. The most natural definition to use is a hybrid of the ‘focal-complement’ definition and the ‘individual-centered’ definitions discussed in (Kerr *et al.* [2004]).

Neighbor altruism: A is an altruist iff,

$$\alpha_i < \beta_i, \quad \text{for } i \in \{0, 1, 2, \dots, n\} \quad (1)$$

$$\alpha_i < \alpha_{i+1}, \quad \text{for } i \in \{0, 1, 2, \dots, n-1\} \quad (2)$$

$$\beta_i < \beta_{i+1}, \quad \text{for } i \in \{0, 1, 2, \dots, n-1\} \quad (3)$$

Both types have higher fitness as a function of the number of A types among their neighbors, and the B type has higher fitness than A in any given neighborhood.

In group-structured models, a *group frequency distribution* is needed in order to make evolutionary predictions. In the neighbor-structured cases, this becomes a *neighborhood frequency distribution*. If individuals are distributed on the space randomly, then the distribution of neighborhoods experienced by both types is the same, and follows a binomial distribution with parameters n and p . Writing the average fitness of the A type as W_A , and that of B type as W_B :

$$W_A = \sum_{i=0}^n \binom{n}{i} p^i (1-p)^{n-i} (\alpha_i) \quad (4)$$

$$W_B = \sum_{i=0}^n \binom{n}{i} p^i (1-p)^{n-i} (\beta_i) \quad (5)$$

After calculating mean fitness in the population \bar{W} as the frequency-weighted average of these fitnesses, $(pW_A + (1-p)W_B)$, the frequency of the A type in the next generation p' can be calculated as pW_A/\bar{W} .

In the case of altruism, however, a simple result follows immediately. From (4) and (5) we have:

$$W_B - W_A = \sum_{i=0}^n \binom{n}{i} p^i (1-p)^{n-i} (\beta_i - \alpha_i) \quad (6)$$

By condition (1) above, $W_B - W_A > 0$, or $W_B > W_A$. So under random neighborhood formation, neighbor altruists have a lower average fitness regardless of their frequency, and thus will go extinct.

In an ‘ephemeral’ neighbor-structured case, of the kind discussed in this paper, a random distribution (or some other distribution) can be restored each generation. In the ‘viscous’ populations that are discussed more often, reproduction has local effects on the distribution of neighborhoods in the next round, and an initially random distribution does not remain random for long. The distribution of neighborhoods co-evolves with the population in a complicated way.

In the remainder of this Appendix, I look at the role of nonrandom distributions in ephemeral neighbor-structured populations.

Nunney ([1985]) represents nonrandom distributions of neighborhoods by using a parameter F , derived from models of inbreeding (see also Hamilton [1975]). If the overall frequencies of A and B are p and q , then p_A and p_B are the frequencies of the A type *experienced* as neighbors by the A type and by the B type, respectively. That is, p_A can be thought of as the frequency of A -type neighbors given the assumption of an A type in focal position. These ‘experienced’ frequencies are calculated by Nunney as follows:

$$p_A = p + (1 - p)F \quad (7)$$

$$p_B = p - pF \quad (8)$$

F ranges between 1 (complete clumping) and 0 (random distribution). Then Nunney treats the distribution of neighborhoods for A as binomial with parameters p_A and n ; the distribution for B is binomial with parameters p_B and n . If these statistics can be assumed to characterize a distribution of neighborhoods, then p_A can be substituted for p in formula (4), and p_B substituted for p in formula (5), to yield average fitnesses of the two types. The ordinary frequency p (not p_A) is then used in the formulas for \bar{W} and p' , to yield a formula for change, $p' = pW_A/\bar{W}$ as before.

However, in many cases it is not straightforward to give a rule for filling the space that will generate a distribution of neighborhoods with these features. In a one-dimensional space with no loop, there is no problem. We can set the first cell randomly (via p) and fill left to right via a first-order Markov process. A transition matrix where P_{ij} is the probability of the next cell being in state j given that the previous cell was in state i can be constructed as follows:

$$\begin{bmatrix} p_A & 1 - p_A \\ p_B & 1 - p_B \end{bmatrix} \quad (9)$$

If all entries in the matrix are positive (which requires $F < 1$), the Markov process will produce limiting global frequencies of p and $(1 - p)$ for the A and B

types. In Case 3 in the main text, $n = 1$ so the resulting formula for change is simple.

$$\begin{aligned}
 p' &= p[p_A \alpha_1 + (1 - p_A) \alpha_0] / \bar{W}, \\
 \bar{W} &= p[p_A \alpha_1 + (1 - p_A) \alpha_0] + (1 - p)[p_B \beta_1 + (1 - p_B) \beta_0]
 \end{aligned}
 \tag{10}$$

Complications arise for dimensions higher than one if we assume that the space is closely packed with individuals as in Figure 2 above. (Nunney's cases are presented differently from this.) Then it can be hard to give a procedure for filling the space that will generate those experienced distributions of neighborhoods for the two types. If we imagine starting with a random choice and using p_A and p_B to choose the state of subsequent cells, we encounter the problem that each cell's state should be constrained by several of its neighbors, not just the one filled previously. A simple way to construct a two-dimensional case that fits something like Figure 2 and has a reasonably simple pattern of correlation might be to suppose that each row begins separately with a random choice and is filled left to right via application of p_A and p_B in a Markov process. Then, with the relevant neighborhood as the 4-member Von Neumann neighborhood, there will be correlation with respect to horizontal neighbors but randomness with respect to vertical neighbors, leading to an overall positive degree of correlation. Another possibility is a trial-and-error process of settling and resettling. An initial filling of the lattice according to p and $(1 - p)$ is followed by a repeated process in which pairs of individuals are randomly chosen and assessed for whether a switch of their places would result in an improved fit to the desired overall distribution of neighborhoods.

Acknowledgements

I am grateful to Peter Koellner, Elisabeth Lloyd, Samir Okasha, Ken Reisman, Susanna Rinard, Michael Weisberg and Rasmus Winther for comments on an earlier draft. Two anonymous referees made helpful criticisms. I am especially grateful to Ben Kerr for detailed comments and assistance.

*Department of Philosophy
Emerson Hall
Harvard University
Cambridge
MA 02138
USA
pgs@fas.harvard.edu*

References

- Alexander, J. [2003]: 'Random Boolean Networks and Evolutionary Game Theory', *Philosophy of Science*, **70**, pp. 1289–304.
- Alexander, J. and Skyrms, B. [1999]: 'Bargaining With Neighbors: Is Justice Contagious?', *Journal of Philosophy*, **96**, pp. 588–98.
- Brandon, R. [1982]: 'The Levels of Selection', in P. Asquith and T. Nickles (eds), *PSA 1982, Volume 1*. East Lansing: Philosophy of Science Association, pp. 315–22.
- Buss, L. [1987]: *The Evolution of Individuality*, Princeton: Princeton University Press.
- Damuth, J. and Heisler, L. [1988]: 'Different Formulations of Multi-Level Selection', *Biology and Philosophy*, **3**, pp. 407–30.
- Frank, S. [1998]: *Foundations of Social Evolution*, Princeton: Princeton University Press.
- Godfrey-Smith, P. and Lewontin, R. C. [1993]: 'The Dimensions of Selection', *Philosophy of Science*, **63**, pp. 373–95.
- Goodnight, C., Schwartz, J. and Stevens, L. [1992]: 'Contextual Analysis of Models of Group Selection, Soft Selection, Hard Selection, and the Evolution of Altruism', *American Naturalist*, **140**, pp. 743–61.
- Hamilton, W. D. [1964]: 'The Genetical Evolution of Social Behavior, I and II', *Journal of Theoretical Biology*, **7**, pp. 1–16 and 17–52.
- Hamilton, W. D. [1975]: 'Innate Social Aptitudes of Man: An Approach from Evolutionary Genetics', reprinted in *Narrow Roads of Gene Land: Collected Papers of W. D. Hamilton, Volume 1*, Oxford: Oxford University Press, 1998, pp. 329–51.
- Hull, D. [1980]: 'Individuality and Selection', *Annual Review of Ecology and Systematics*, **11**, pp. 311–32.
- Kerr, B. and Godfrey-Smith, P. [2002]: 'Individualist and Multi-level Perspectives on Selection in Structured Populations', *Biology and Philosophy*, **17**, pp. 477–517
- Kerr, B., Godfrey-Smith, P. and Feldman, M. [2004]: 'What is Altruism?', *Trends in Ecology and Evolution*, **19**, pp. 135–40.
- Leiberman, E., Hauert, C. and Nowak, M. [2005]: 'Evolutionary Dynamics on Graphs', *Nature*, **433**, pp. 312–6.
- Lloyd, E. [2001]: 'Units and Levels of Selection: An Anatomy of the Units of Selection Debates', in R. Singh, C. Krimbas, D. Paul and J. Beatty (eds), *Thinking About Evolution: Historical, Philosophical and Political Perspectives*, Cambridge: Cambridge University Press, pp. 267–91.
- Lloyd, E. [2005]: 'Why the Gene Will Not Return', *Philosophy of Science*, **72**, pp. 287–310.
- Maynard Smith, J. [1964]: 'Group Selection and Kin Selection', *Nature*, **201**, pp. 1145–7.
- Maynard Smith, J. [1976]: 'Group Selection', *Quarterly Review of Biology*, **51**, pp. 277–83.
- Maynard Smith, J. [1987]: 'How to Model Evolution', in J. Dupré (ed.) *The Latest on the Best: Essays on Evolution and Optimality*, Cambridge, MA: MIT Press, pp. 119–31.
- Maynard Smith, J. [2002]: 'Commentary on Kerr and Godfrey-Smith', *Biology and Philosophy*, **17**, pp. 523–7.
- Maynard Smith, J. and Szathmari, E. [1995]: *The Major Transitions in Evolution*, Oxford: Oxford University Press.

- Michod, R. [1982]: 'The Theory of Kin Selection', *Annual Review of Systematics*, **13**, pp. 23–55.
- Michod, R. [1999]: *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*, Princeton: Princeton University Press.
- Mitteldorf, J. and Wilson, D. S. [2000]: 'Population Viscosity and the Evolution of Altruism', *Journal of Theoretical Biology*, **204**, pp. 481–96.
- Nowak and May [1992]: 'Evolutionary Games and Spatial Chaos', *Nature*, **359**, pp. 826–29.
- Nunney, L. [1985]: 'Group Selection, Altruism, and Structured-Deme Models', *American Naturalist*, **126**, pp. 212–30.
- Okasha, S. [2004]: 'Multi-Level Selection, Covariance and Contextual Analysis', *British Journal for the Philosophy of Science*, **55**, pp. 481–504.
- Okasha, S. [2005]: 'Altruism, Group Selection and Correlated Interaction', *British Journal for the Philosophy of Science*, **56**, pp. 703–25.
- Okasha, S. [2006]: *Evolution and the Level of Selection*, Oxford: Oxford University Press.
- Price, G. [1972]: 'Extension of Covariance Selection Mathematics', *Annals of Human Genetics*, **35**, pp. 485–90.
- Rice, S. [2004]: *Evolutionary Theory: Mathematical and Conceptual Foundations*, Sunderland: Sinauer.
- Rinard, S. [unpublished]: 'Sober and Wilson on Group Selection'.
- Sarkar, S. [2008]: 'A Note on Frequency-Dependence and the Levels/Units of Selection', *Biology and Philosophy*, **23**, pp. 217–28.
- Skyrms, B. [1994]: 'Darwin Meets the Logic of Decision: Correlation in Evolutionary Game Theory', *Philosophy of Science*, **61**, pp. 503–28.
- Skyrms, B. [2004]: *The Stag Hunt and the Evolution of Social Structure*, Cambridge: Cambridge University Press.
- Sober, E. [1984]: *The Nature of Selection*, Cambridge, MA: MIT Press.
- Sober, E. [1992]: 'The Evolution of Altruism: Correlation, Cost and Benefit', *Biology and Philosophy*, **7**, pp. 177–88.
- Sober, E. and Lewontin, R. C. [1982]: 'Cause, Artifact, and Genic Selection', *Philosophy of Science*, **47**, pp. 157–80.
- Sober, E. and Wilson, D. S. [1998]: *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Cambridge, MA: Harvard University Press.
- Sterelny, K. [1996]: 'The Return of the Group', *Philosophy of Science*, **63**, pp. 562–84.
- Uyenoyama, M. and Feldman, M. [1980]: 'Theories of Kin and Group Selection: A Population Genetics Perspective', *Theoretical Population Biology*, **17**, pp. 380–414.
- Wade, M. [1985]: 'Soft Selection, Hard selection, Kin Selection, and Group Selection', *American Naturalist*, **125**, pp. 61–73.
- Wallace, B. [1968]: 'Polymorphism, Population Size, and Genetic Load', in R. C. Lewontin (ed.), *Population Biology and Evolution*, Syracuse: Syracuse University Press, pp. 87–108.
- Waters, K. [2005]: 'Why Genic and Multi-Level Theories are Here to Stay', *Philosophy of Science*, **72**, pp. 311–33.

- Werfel, J. and Bar-Yam, Y. [2004]: 'The Evolution of Reproductive Restraint Through Social Communication', *Proceedings of the National Academy of Sciences, USA*, **101**, pp. 11019–24.
- Williams, G. C. [1966]: *Adaptation and Natural Selection*, Berkeley: University of California Press.
- Wilson, D. S. [1975]: 'A Theory of Group Selection', *Proceedings of the National Academy of Sciences, USA*, **72**, pp. 143–6.
- Wilson, D. S. [1977]: 'Structured Demes and the Evolution of Group-Advantageous Traits', *American Naturalist*, **111**, pp. 157–85.
- Wilson, D. S. [1980]: *The Natural Selection of Populations and Communities*, Menlo Park: Benjamin/Cummins.
- Wilson, D. S. and Dugatkin, L. [1997]: 'Group Selection and Assortative Interactions', *American Naturalist*, **149**, pp. 336–51.
- Wilson, D. S., Pollock, G. and Dugatkin, L. [1992]: 'Can Altruism Evolve in Purely Viscous Populations?', *Evolutionary Ecology*, **6**, pp. 331–41.