

Why Won't the Group Selection Controversy Go Away?

Samir Okasha

ABSTRACT

The group selection controversy is about whether natural selection ever operates at the level of groups, rather than at the level of individual organisms. Traditionally, group selection has been invoked to explain the existence of altruistic behaviour in nature. However, most contemporary evolutionary biologists are highly sceptical of the hypothesis of group selection, which they regard as biologically implausible and not needed to explain the evolution of altruism anyway. But in their recent book, Elliot Sober and David Sloan Wilson [1998] argue that the widespread opposition to group selection is founded on conceptual confusion. The theories that have been propounded as alternatives to group selection are actually group selection in disguise, they maintain. I examine their arguments for this claim, and John Maynard Smith's arguments against it. I argue that Sober and Wilson arrive at a correct position by faulty reasoning. In the final section, I examine the issue of how to apply the principle of natural selection at different levels of the biological hierarchy, which underlies the dispute between Sober and Wilson and Maynard Smith.

- 1 *Introduction*
 - 2 *Sober and Wilson's defence of group selection*
 - 3 *Sober versus Maynard Smith [1987]*
 - 4 *Sober and Wilson versus Maynard Smith [1998]*
 - 5 *Why the concept of group heritability is irrelevant*
 - 6 *On hierarchical application of the principle of natural selection*
 - 7 *Conclusion*
-

1 Introduction

Orthodox Darwinian theory treats the individual organism as the basic 'unit of selection'. To see what this means, consider a typical Darwinian explanation, of why the average running speed in an antelope population has increased over time, for example. In the ancestral population, individual antelopes varied with respect to running speed; faster antelopes were better able to avoid predators, so tended to leave more offspring; and the offspring of fast antelopes tended to be fast runners themselves. Thus over time we

should expect the average running speed in the population to increase, *ceteris paribus*. In explanations of this sort, the change in frequency of the phenotypic trait of interest—running fast—is explained by the advantage that the trait confers on individual organisms that possess it. That is what it means to say that the ‘unit of selection’ is the individual.¹ The group selection controversy is about whether *groups* of organisms can sometimes be the units of selection. Might certain traits evolve and persist because of the advantage they confer on *groups* of organisms in their struggle for survival over other groups, rather than on individual organisms themselves?

The origins of the group selection debate lie in Darwin’s discussion of altruistic behaviour among humans in *The Descent of Man* (Darwin [1871]). The existence of altruism poses a *prima facie* problem for the theory of natural selection. For behaving altruistically *lowers* an individual’s fitness relative to that of selfish individuals, and should thus be disfavoured by natural selection. As Darwin put it: ‘he who was ready to sacrifice his life, as many a savage has been, rather than betray his comrades, would often leave no offspring to inherit his noble nature’ ([1871], p. 163). So how can altruism evolve? Darwin’s answer was to appeal to selection at the level of the *group*. Though altruistic individuals do worse than selfish ones, it is quite possible that *groups* with a high proportion of altruists might do better than *groups* with a lower proportion. Darwin wrote: ‘a tribe including many members who . . . were always ready to give aid to each other and sacrifice themselves for the common good, would be victorious over most other tribes; and this would be natural selection’ ([1871], p. 166). Darwin’s suggestion is that group selection for altruism may have outweighed the individual selection against, thus explaining the existence of altruism in nature today.

The problem of altruism remains at the conceptual centre of the group selection controversy today. A very brief history of the contemporary controversy looks like this. During the 1950s and 1960s, biologists routinely interpreted much animal behaviour as an adaptation designed to benefit the group or the species. This view of evolution received its clearest formulation in the work of the Scottish biologist Wynne-Edwards ([1962]). He argued, for example, that individual organisms often deliberately forego reproduction when the size of their group is getting too high, in order to avoid over-exploiting the group’s limited resources. Such behaviour had evolved because groups *not* containing individuals who practised reproductive restraint had gone extinct, and their habitats were colonized by offshoots from groups that did contain such individuals. Wynne-Edwards saw no objection—and no alternative—to invoking between-group selection to explain altruistic

¹ Throughout I use the expression ‘individual’ to mean ‘individual organism’. This is not supposed to pre-judge the issue of whether certain higher-level entities, such as species, count as individuals in a generic sense, as many philosophers of biology maintain.

behaviour in nature. But all this changed in the mid-1960s, thanks to the powerful critiques of group selection by George Williams ([1966]) and John Maynard Smith ([1964], [1976]). These authors argued that group selection is biologically unlikely. Since the lifetime of a group is much longer than that of an individual, individual selection will invariably be a stronger force in nature than group selection. This view was supported by a number of mathematical models, which suggested that group selection would only have evolutionarily significant effects for a very restricted set of parameter values (Levins [1970]; Boorman and Levitt [1973]; Levin and Kilmer [1974]). Furthermore, Maynard Smith and Williams argued that group selection was not needed to explain the evolution of altruistic behaviour anyway—there were alternative explanations. Not only was the hypothesis of group selection implausible, it was also explanatorily superfluous, they maintained.

Chief among the alternative explanations of altruism was the idea of *kin selection*. Kin selection theory—also known as ‘inclusive fitness theory’—developed out of W. D. Hamilton’s seminal work on the evolution of social behaviour (Hamilton [1964a], [1964b]). Hamilton developed an insight that had been vaguely appreciated since Darwin’s day: organisms that invest energy in caring for their offspring will have an obvious selective advantage over those that do not. ‘Caring for offspring’ is thus an altruistic behaviour which can evolve by ordinary, individual-level natural selection. Hamilton’s theory generalized this insight in a precise mathematical form. He showed that a gene which codes for behaviour that is costly to an individual but benefits his relatives, e.g. sharing food with siblings, can increase in frequency by natural selection—because the individual’s relatives are likely to carry copies of the gene in question themselves. Altruism can evolve, Hamilton concluded, so long as the cost incurred by the altruist is offset by a sufficient amount of benefit to sufficiently closely related relatives.² Hamilton’s work was initially aimed at explaining co-operative behaviour in social insect castes, but his theory quickly came to be applied to altruistic behaviour in general. Kin selection seemed to provide a way of explaining the evolution of altruism *without* positing a process of between-group selection, and was widely seen as a superior alternative to Wynne-Edwards-style group selection.

Kin selection was not the only alternative explanation of altruism that emerged out of discontent with traditional group selection. Another idea, better able to explain altruism among non-relatives, was the ‘evolutionary game theory’ of Maynard Smith ([1982]); closely related ideas were developed

² This condition for the evolution of altruism is known as *Hamilton’s Rule*. In its simplest form, the rule states that an allele which codes for altruistic behaviour will spread so long as $b/c > 1/r$, where c is the cost incurred by the altruist (the donor), b the total benefit enjoyed by recipients of the altruism, and r the coefficient of relationship between donor and recipient. See Grafen ([1984]) for useful discussion.

by Trivers ([1971]) and Axelrod and Hamilton ([1981]). Evolutionary game theory is a way of modelling the fitness consequences of social interactions between pairs of individual organisms. In the simplest case, each organism is either selfish or altruistic. When an altruist interacts with a selfish individual, the latter does best; when two altruists interact, both do very well; when two selfish individuals interact, both do very badly. Payoffs from the three possible types of interaction can be represented in a 2×2 matrix. Once the payoff values have been specified, and once the frequency of each type of interaction is fixed, it can be determined whether altruism or selfishness will evolve. Intuitively, altruism can evolve so long as altruist–altruist interactions are sufficiently common, and sufficiently favourable, to offset the advantage that selfish individuals have in selfish–altruist interactions. As with kin selection theory, evolutionary game theory seemed to provide a way of explaining the evolution of altruism which did *not* involve a process of between-group selection, and which was not susceptible to the criticisms levelled against traditional group selection theory.

Due to the combined influence of Williams', Maynard Smith's and Hamilton's work, the idea of group selection fell into widespread disrepute in mainstream evolutionary biology, where it remains today. As D. S. Wilson says, among many contemporary biologists 'group selection is treated as such heresy that the only thing to learn about it is "Just say No"' ([1997], p. S2). But the controversy has refused to die completely. Despite the best efforts of its opponents, group selectionist thinking has by no means been expunged entirely from evolutionary biology, and remains popular in some quarters. In this paper I examine some recent salvos in the group selection controversy. I focus on a recent book-length defence of group selection by Elliot Sober and David Sloan Wilson ([1998]), and the critical reaction to their book by John Maynard Smith ([1998]).

Firstly, it is worth explaining briefly why the group selection debate should be of concern to a philosopher. For the issue may *look* straightforwardly empirical—surely it is about whether a certain causal process, called 'group selection', has or has not played an important role in evolutionary history? Finding out the answer may be difficult, but how can a philosopher help? In fact matters are not quite so simple. As will become clear below, the group selection debate actually involves a curious blend of empirical and conceptual issues, which makes it ideally suited to—and much in need of—philosophical clarification. Here I add my name to the list of those philosophers who have attempted to provide the requisite clarification.³

³ For previous attempts, see Cassidy ([1978]), Sober ([1984]), Sterelny ([1996]), and Sterelny and Griffiths ([1999]), Ch. 8.

2 Sober and Wilson's defence of group selection

Sober and Wilson have been enthusiastic supporters of group selection for a number of years.⁴ In their book *Unto Others* ([1998]), they attack the widespread consensus that group selection is an unimportant evolutionary force. Their case is twofold. Firstly, they argue that the standard arguments against group selection are not as powerful as most people believe. Secondly and more importantly, they argue that the alternative explanations of altruistic behaviour that emerged in the 1960s and 1970s, such as kin selection, are not actually alternatives to group selection at all! Kin selection theory and evolutionary game theory are actually *versions* of group selection, according to Sober and Wilson. They write: 'the theories that have been celebrated as alternatives to group selection are nothing of the sort. They are different ways of viewing evolution in multi-group populations . . . However, the theories are formulated in a way that obscures the role of group selection' ([1998] p. 57). In other words, those very biologists who regard themselves as vehement *opponents* of traditional group selection are actually group selectionists in disguise; for group selection is implicated in the alternative explanations of altruism that these biologists promote, Sober and Wilson contend.

Though this claim is *prima facie* startling, it is not *quite* as radical as it appears at first sight. For as Sober and Wilson note, a small minority of evolutionary theorists have come to the conclusion that group selection theory was prematurely rejected in the 1970s, and that kin selection, in particular, can be viewed as a type of group selection. Indeed, they cite some intriguing evidence that W. D. Hamilton himself endorsed this position. Sober and Wilson argue that Hamilton started out as an avid opponent of group selection, but later came round to the view that his own explanations of altruism actually involved an element of group selection, just as Sober and Wilson maintain. Hamilton (personal communication) has confirmed to me that Sober and Wilson's account of his changing attitude towards group selection is essentially accurate. But none the less, Sober and Wilson's position *is* radical. The majority of evolutionary biologists, and virtually all those non-biologists with a working knowledge of evolutionary theory, would be astonished by Sober and Wilson's thesis that the theories usually regarded as alternatives to group selection are actually group selection in disguise. The thesis therefore merits close attention.

To understand how Sober and Wilson arrive at their position, we need to look briefly at Wilson's so-called 'intra-demic' selection model, also known as

⁴ Wilson ([1975], [1977], [1980]) and Sober and Wilson ([1994]) all treat group selection in a manner consonant with their most recent effort.

the ‘trait group’ selection model.⁵ This model was first put forward by Wilson ([1975]), and forms the basis of his and Sober’s case. The model assumes a single, freely mixing population with the following structure. At a certain stage of their life cycle, individual organisms come together in temporary groupings (known as ‘trait groups’), where they interact with one another, and certain behaviours are expressed. For simplicity, assume that the expressed behaviour is either altruistic or selfish, and is under full genetic control. After interacting, the individuals disperse back into the global population and mate randomly; the cycle repeats in the following generation (see Figure 1). Thus for example, many organisms spend their juvenile stage in small sibling groups, before blending into the global population as adults and mating. Altruistic behaviour can evolve, on this model, so long as there is a statistical tendency for altruists to find themselves grouped with other altruists in the pre-dispersal phase. If groupings are formed at random, then altruism is eventually eliminated from the population, for *within* each grouping selfish individuals are fitter than their altruistic colleagues (by definition). But if altruists tend to find themselves grouped with other altruists, and thus tend to be the recipients of each other’s help, then altruism can evolve. When this happens, the fitness disadvantage that altruists face *within* groups is offset by the fact that groups that contain many altruists are fitter than groups that do not. Non-random assortment is the *sine qua non* of the evolution of altruism, on Wilson’s model.

A simple numerical example can help to make this clear (adapted from Sober and Wilson ([1998]) pp. 24–5). Imagine a population containing 300 organisms, half of whom are altruists (A), the other half selfish (S). For simplicity, assume that reproduction is asexual and that like always begets like—the offspring of selfish organisms are selfish, of altruistic organisms altruistic. The population is sub-divided into three temporary groupings of 100, in which fitness-affecting interactions take place. After interacting, the groups dissolve and the organisms blend back into the global population, mix freely and reproduce. For simplicity, assume that organisms die as soon as their reproductive potential is exhausted. Group one contains 80 A and 20 S, group two 60 A and 40 S, and group three 10A and 90S. Behaving altruistically reduces an organism’s fitness, but boosts the fitness of all other organisms in the group. The fitness of any organism (i.e. the expected number of offspring it will leave) thus depends on (a) whether it is altruistic or selfish, and (b) what group it is in. Fitness values are determined as follows. Each organism has a ‘baseline’ fitness of 10. Altruists incur a cost of 1, but raise the fitness of those in their group by a total of 5, shared equally among the other

⁵ My reconstruction of Sober and Wilson’s argument departs somewhat from the order of presentation in their book. This is purely for ease of exposition.

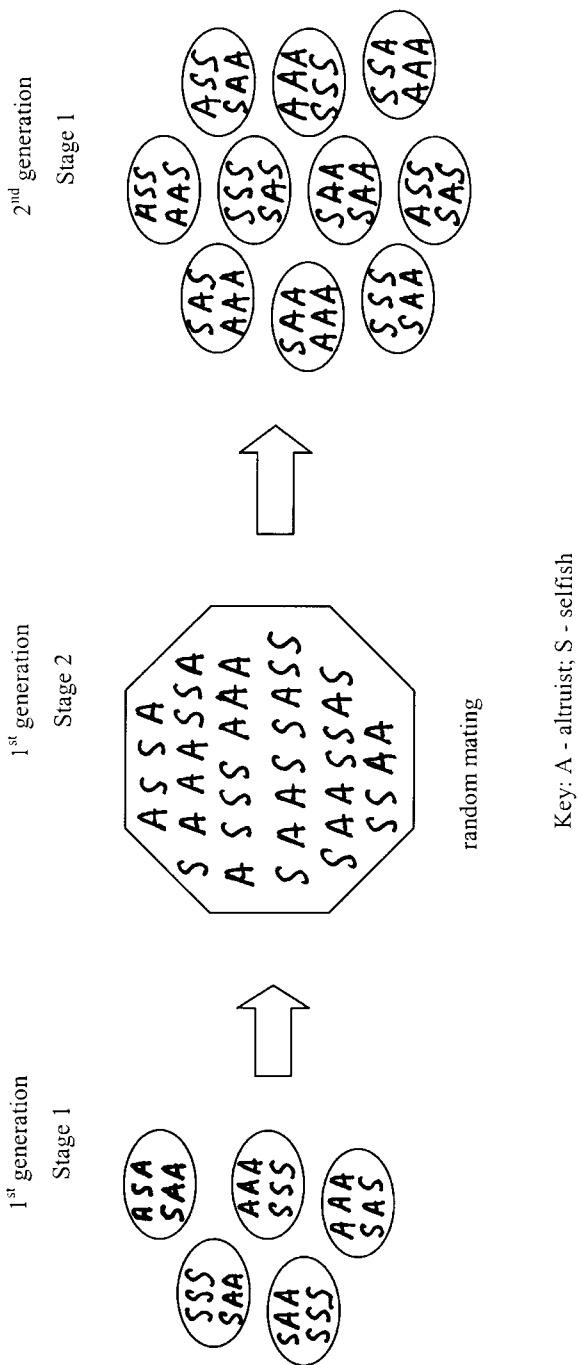


Figure 1: Population Structure in the Intra-Demic Selection Model

99 organisms in the group. So within any group, the fitness of a selfish organism is $10 + 5(x/99)$, while the fitness of an altruist is $10 - 1 + 5(x-1/99)$, where x is the number of altruists in the group.⁶ Actual fitness values are shown in Table 1.

Within each group, selfish individuals are fitter than altruistic ones (i.e. $W_s > W_a$ for each group). But the fitness of an individual depends on which group it is in. Fitnesses are highest in group one, for it contains the highest proportion of altruists, and lowest in group three, for it contains the lowest. So the total number of offspring contributed by members of group one to the global population is the highest, while that of group three is the lowest (i.e. n' is highest for group 1, lowest for group 3). To calculate the relative frequency of altruism and selfishness in the second generation, we multiply the number of altruists in each group by the fitness of altruists in that group, do the same for selfish individuals, then add the results. As Figure 1 indicates, in the global population altruism actually *increases* in frequency, from 0.5 to 0.51. So after selection, the proportion of altruists in the global population is higher than it was before, despite the fact that *within* each group altruists are lower in fitness than their selfish counterparts. This paradoxical-sounding result arises because altruists are concentrated in fitter groups, i.e. groups that contribute more progeny to subsequent generations. The positive correlation between group fitness and frequency of altruists is responsible for the overall increase in altruism, overriding the selective disadvantage that altruists face within each group. In subsequent generations altruism will continue to increase in frequency, so long as altruists continue to find themselves grouped together in sufficient number.

Sober and Wilson insist that intra-demic selection *is* group selection, even though the ‘groups’ in Wilson’s model do not need to be reproductively isolated, and only last for a fraction of the lifetime of individual organisms.⁷ Indeed, Wilson’s ‘groups’ do not even need to be spatially discrete: any collection of organisms whose interactions affect each others’ fitness qualifies as a group, on his view. In traditional Wynne–Edwards style group selection, sometimes called ‘*inter-demic*’ selection, groups *were* reproductively isolated, spatially discrete, and lasted for many generations. These differences are irrelevant, Sober and Wilson argue—for the causal processes at work in intra-demic and traditional group selection are identical. In both cases, altruists are

⁶ The term $5(x-1/99)$ in the expression for the fitness of altruists, as opposed to $5(x/99)$ in the expression for selfish organisms, reflects the fact that selfish organisms receive benefit from all altruists in their group, while altruists receive benefit from all altruists in their group except themselves.

⁷ The issue of reproductive isolation obviously does not arise in the numerical example above, where reproduction was assumed to be asexual for simplicity. But in more realistic intra-demic models that do not assume asexual reproduction, e.g. Wilson ([1975]), the groups are not reproductively isolated (hence the expression ‘*intra-demic*’).

Table 1. The evolution of altruism in a multi-group population.

	Group 1	Group 2	Group 3
n	100	100	100
P _a	0.8	0.6	0.1
T _a	100 * (0.8) = 80	100 * (0.6) = 60	100 * (0.1) = 10
W _a	10 - 1 + 5(79/99) = 12.99	10 - 1 + 5(59/99) = 11.98	10 - 1 + 5(9/99) = 9.45
W _s	10 + 5(80/99) = 14.04	10 + 5(60/99) = 13.03	10 + 5(10/99) = 10.51
n'	(12.99*80) + (14.04*20) = 1320	(11.98*60) + (13.03*40) = 1240	(9.45*10) + (10.51*90) = 1040
Global Population			
N	100 + 100 + 100 = 300		
P _A	(0.8(100) + 0.6(100) + 0.1(100)) / 300 = 0.5		
T _A	300 * 0.5 = 150		
N'	1320 + 1240 + 1040 = 3600		
T _A '	(12.99*80) + (11.98*60) + (9.45*10) = 1852.5		
P _A '	1852.5 / 3600 = 0.51		

Key: n = number of 1st generation organisms per group; n' = number of offspring contributed by members of a group to global population; P_a = frequency of altruists per group; T_a = number of altruists per group; W_a = fitness of altruistic types; W_s = fitness of selfish types; N = total number of 1st generation organisms; N' = total number of 2nd generation organisms; P_A = global frequency of altruists in 1st generation; T_A = total number of altruists in 1st generation; P_A' = global frequency of altruists in 2nd generation; T_A' = total number of altruists in 2nd generation

less fit than selfish individuals *within* each group, but groups containing many altruists are fitter than groups containing fewer. So in both cases, individual selection operates on fitness differences *within* groups and favours selfishness, but group selection operates on fitness differences *between* groups, and favours altruism. Intra-demic selection has as good a right to be considered group selection as the process Wynne-Edwards envisaged, Sober and Wilson insist.

Furthermore, Sober and Wilson argue that kin selection and evolutionary game theory are simply special cases of intra-demic selection. This point is easily seen in the case of kin selection, where the entities corresponding to the groups of the intra-demic model are kin groups. The evolution of altruism by kin selection proceeds exactly as in the example above: the individual disadvantage of behaving altruistically is offset by the fact that altruists tend to be the recipients of each other's help. The division of the population into kin groups is simply an efficient way of getting the positive assortment of altruists with one another that the evolution of altruism requires.⁸

⁸ Sober and Wilson are obviously right to stress that what matters to the evolution of altruism is that the recipients of altruistic behaviour be altruists themselves, not that they be relatives of the donor. However, I suspect that most biologists would regard kinship as the only biologically realistic way of securing positive assortment of altruists, generation after generation. Sober and Wilson disagree—they apparently regard positive assortment of unrelated altruists as not just conceivable but frequent in nature. See Grafen ([1984]) and Maynard Smith ([1976], [1987b]) for discussion of this important empirical issue.

Evolutionary game theory can also be conceptualized as intra-demic selection, Sober and Wilson argue, by regarding the *interacting pairs* as the groups. Again, altruists face a disadvantage *within* groups (i.e. they do worse in selfish–altruist interactions), but groups of altruists (i.e. altruist–altruist pairs) are fitter than groups of the other two types. So although kin selection and evolutionary game theory are not usually described in group selectionist terms, in fact they presume a population structure identical to that of Wilson’s intra-demic model; and intra-demic selection *is* group selection, according to Sober and Wilson. Their case for the thesis that the supposed alternatives to group selection are actually group selection in disguise thus comprises two distinct claims: (a) the supposed alternatives to group selection are cases of intra-demic selection, and (b) intra-demic selection is itself group selection.

Of these two claims, (b) is the more controversial. For it is really just a mathematical fact that kin selection and evolutionary game theory are instances of intra-demic selection (see Charnov ([1977]) or Wilson ([1977]) for proof of this point in relation to kin selection; see Sober and Wilson ([1994]) for informal proof in relation to evolutionary game theory). Admittedly, there is no overriding reason to re-describe either kin selection or evolutionary game theory within the framework of intra-demic selection, as Sober and Wilson admit; but this does not alter the fact that such re-description is possible. However, the claim that intra-demic selection is itself a type of group selection *is* controversial, and it is here that the majority of evolutionary biologists would quarrel with Sober and Wilson’s reasoning. I focus on this claim in what follows.

In a review of their book in *Nature*, John Maynard Smith claims that the issue at stake between Sober and Wilson and orthodox evolutionary biologists like himself is ‘semantic’, for it ‘could not be settled by observation’ ([1998], p. 639). Sober and Wilson have simply chosen to use the expression ‘group selection’ in a non-standard and highly misleading way, he claims; so their departure from orthodoxy is purely terminological.

Maynard Smith is certainly right that the issue cannot be settled by observation, and he is right, I think, that there is a process which Sober and Wilson call ‘group selection’ and which other people do not. But the inference from ‘no observations can decide the issue’ to ‘the issue is semantic’ seems to me a bit too fast. I think a deeper diagnosis of what is going on is possible. As all philosophers know, any two objects are similar in some respects but dissimilar in others. In devising concepts with which to describe the world, we must therefore select some similarity relations as salient, and ignore others. Which similarity relations we select may depend on our background beliefs, our theoretical interests, and so on. Thus a dispute about whether two objects are of the same type is a dispute about which concepts are most appropriate for describing our world, and that is something that may lend itself to rational discussion, even though it is not decidable by observation.

Something like this applies to the group selection controversy, in my view. Here the objects that need classifying are selection processes, and the relevant categories are 'individual selection', 'kin selection', 'group selection' etc. All parties agree that the evolution of running speed among antelope described earlier is a case of individual selection; and all parties agree that the process envisaged by Wynne-Edwards—differential extinction and colonization of reproductively isolated groups—is a case of group selection. But there are other processes, such as those described by Wilson's intra-demic selection model, which are similar to traditional group selection in some respects but not others; and different theorists disagree about the *relevance* of these similarities, for a variety of reasons. In order to apply this diagnosis to the current dispute between Sober and Wilson and Maynard Smith, it is first necessary to examine an earlier clash between Sober and Maynard Smith, on the same issue.

3 Sober versus Maynard Smith [1987]

In an exchange of views with Sober in 1987, Maynard Smith argued that Wilson's intra-demic selection model did *not* count as group selection, for a very simple conceptual reason. Natural selection occurs whenever we have a set of objects that satisfy three conditions. Firstly, the objects vary with respect to possession of a given trait; secondly, this variation is heritable, i.e. passed on from parents to offspring; and thirdly, possession of the trait influences the expected number of offspring an object leaves, i.e. the object's fitness. In short, natural selection requires 'heritable variation in fitness'.⁹ In individual selection, the objects that vary in fitness are individual organisms; they vary in fitness because their phenotypic traits are different; since organisms tends to inherit their parents' phenotypes, the variation in fitness is heritable. For group selection to occur, Maynard Smith reasoned, we need *groups* of organisms that exhibit heritable variation in fitness. And that is what we have, in the process envisaged by Wynne-Edwards. Some groups 'reproduce' by sending out offspring groups to colonize the habitats of groups that have 'died' or gone extinct; this variance in group fitness is caused by differences in group-level traits (e.g. 'frequency of altruists'); since offspring groups inherit the traits of the parental groups they have come from, the variance in fitness is heritable. But, Maynard Smith argued, in Wilson's model of intra-demic selection, where individuals interact with one another in temporary groupings for a short time, then blend into the global population and mate, there *is* no group reproduction—and thus no group heritability. So

⁹ This formulation of the conditions required for evolution by natural selection was first given by Lewontin ([1970]), and has become standard. Maynard Smith expresses the point slightly differently by saying that natural selection requires entities that exhibit 'multiplication, heredity and variation' ([1987a], p. 121).

the group *cannot* be the unit of selection. Intra-demic selection should thus be conceptualized as ordinary individual selection, Maynard Smith argued, but with the fitness of any individual organism partially determined by the group to which it belongs.¹⁰

In reply to Maynard Smith, Sober wrote: ‘I agree that the process Wynne–Edwards postulated involves groups that exhibit heredity. David Wilson’s . . . groups require no such thing’ ([1987], p. 133). But this difference is irrelevant, Sober argued, because group heritability is not needed for group selection to occur.¹¹ He wrote: ‘for natural selection to produce evolution, heritability of some sort is essential. But for *group* selection to cause evolution it is not essential that the heritability be *group* heritability’ (*Ibid.*, p. 136, Sober’s emphasis). So Sober is saying that the process Wilson describes *is* group selection, even though it is individual organisms that stand in parent–offspring relations in that process, not groups. By Sober’s lights, group selection occurs whenever the fitness of an individual organism depends on properties of the group to which it belongs, even if the group is a temporary coalition of organisms; while by Maynard Smith’s lights, group selection requires the differential reproduction and extinction of groups themselves. Maynard Smith reacted to Sober’s claim that group heritability is not a necessary condition for group selection by saying: ‘Sober does not understand what the group selection debate is about’ ([1987b], p. 147).

It looks as if Maynard Smith is indeed on stronger ground here. Maynard Smith’s concept of group selection seems to arise very naturally from applying the familiar idea of natural selection to a different level of the biological hierarchy—to groups rather than individuals. Whereas Sober’s concept seems to introduce a curious asymmetry: individual selection requires individual-level heritability, but group selection does *not* require group-level heritability—though it may be there anyway. On the face of it, Sober’s concept seems to unite selection processes whose underlying causal structures are very different, and which therefore surely belong apart. So my *first* reaction is to say that Maynard Smith has the better of the 1987 exchange; though for reasons that will emerge, that is not my last reaction.

4 Sober and Wilson versus Maynard Smith [1998]

Moving forward eleven years to the recent clash between the same authors, matters are a bit different. Maynard Smith’s position remains the same, but Sober has undergone a dramatic, and unacknowledged, change of mind. For

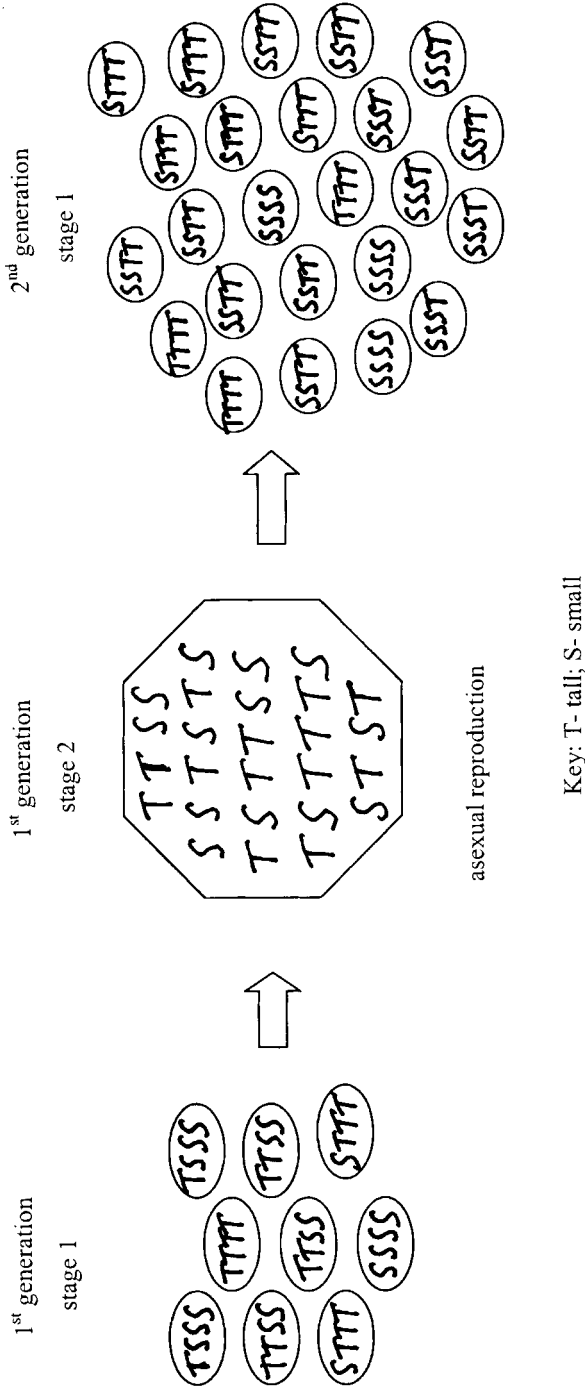
¹⁰ In an earlier paper, Maynard Smith similarly wrote ‘the extinction of some groups, and the ‘reproduction’ of others, are essential features of evolution by group selection’ ([1976], p. 279). Wade ([1978]) argues that extinction is *not* a necessary condition of group selection.

¹¹ Sober uses the phrases ‘group heritability’ and ‘group heredity’ interchangeably.

Sober and Wilson now say that group selection *does* require group heritability—the very point Sober denied in 1987. They argue at length that natural selection at any given level of the biological hierarchy requires heritable variation in fitness at that level—the very point Maynard Smith had insisted on in 1987. But they still maintain that intra-demic selection is group selection, not individual selection. For they now argue that there *is* actually group reproduction and group heritability in Wilson's intra-demic selection model—the very opposite of what Sober said in 1987.


Recall how the intra-demic model works—individuals assort in groups for part of their life cycle, interact, then disperse into the global population and mate randomly; their offspring settle into groups themselves, and the cycle repeats. Sober and Wilson say that the concept of group reproduction *can* apply here—via the following convention. We treat a group in the first generation as the 'parent' of a group in the second generation, just if some individual in the first generation group is the parent of some individual in the second generation group ([1998] p. 111). So most offspring groups will have multiple parents. Given this way of identifying parent–offspring relations between groups, we can ask whether a given group-level trait tends to be passed on from parental groups to offspring groups, i.e. whether it is heritable. So the concept of group heritability *can* apply, even though groups are transient entities whose duration is shorter than that of their component organisms.

Sober and Wilson illustrate their new-found belief in group heritability with a simplified example ([1998] pp. 110–2). Consider a population containing two types of organisms, tall (T) and short (S), in equal proportion. Reproduction is asexual, and like begets like: the offspring of tall individuals are tall, of short individuals short. Organisms spend their juvenile stages in groups of four, then blend into the global population as adults and reproduce; their offspring settle at random into groups of four, and the cycle repeats (Figure 2(a)). There are five possible types of group: {TTTT}, {TTTS}, {TTSS}, {TSSS} and {SSSS}, each of which differs with respect to the group-level trait 'average height'. To assess the heritability of this trait, we pick a given group, of type {TTTT} for example, and find its 'offspring' groups in the second generation, by applying the convention on group reproduction. Since 2nd generation organisms settle at random, the offspring of our 1st generation {TTTT} group will include groups of types {TTTT}, {TTTS}, {TTSS} and {TSSS}, but not {SSSS} obviously. We then calculate the average height of each of the offspring groups. We do the same for groups of each of the other four types. We can now investigate whether or not there is a correlation between parental groups and offspring groups with respect to the trait 'average height'. Do taller parental groups tend to produce taller than average offspring groups or not? The answer is 'yes' (intuitively, because



Sober and Wilson's convention: a 1st generation group is the 'parent' of a 2nd generation group just if some individual in the former is the parent of some individual in the latter.

Figure 2(a): Group ancestry despite blending of groups in each generation



Key: S - small; T - tall

Figure 2(b): Heritability of the group-level trait ‘average height’.
(adapted from Sober and Wilson [1998] p. 212).

Notes: (i) numerical units are arbitrary; (ii) slope of the correlation function is determined by the following assumptions: (a) in both generations, the population contains S and T individuals in equal proportion, and (b) groupings are formed at random, so in both generations average group height is normally distributed (i.e. {SSSS} and {TTTT} groups are rarest, while {SSTT} groups are commonest).

{TTTT} groups are tallest, and none of their offspring groups will be of the smallest {SSSS} type.) The group-level trait ‘average height’ is therefore heritable. This is reflected by the positive slope of the graph in Figure 2(b), which plots the height of parental groups against the average height of their offspring groups.¹²

The population structure in the example above is of course identical to that of the intra-demic selection model. Sober and Wilson therefore conclude that the concept of group heritability does apply in the intra-demic selection model, *contra* Maynard Smith, and *contra* Sober’s earlier view. Maynard Smith’s insistence that group selection requires group heritability can be

¹² It may look as if *every* group-level trait will automatically be heritable, given Sober and Wilson’s convention for determining group ancestry. If so, this would mark a serious disanalogy with individual selection, for the heritability of traits at the individual level is obviously a thoroughly contingent matter. But in fact, Sober and Wilson’s convention does not have this consequence. Only group-level traits that are an *additive function* of individual-level traits, such as average height, will be automatically heritable, given their convention. Group-level traits that are not additively determined by individual-level traits may or may not be heritable.

Table 2

	Is group heritability necessary for group selection?	Is there group heritability in intra-demic model?	Is intra-demic selection a type of group selection?
Maynard Smith	yes	no	no
Sober [1987]	no	no	yes
Sober and Wilson [1998]	yes	yes	yes

respected, Sober and Wilson now think, compatibly with regarding intra-demic selection as a type of group selection.

Table 2 summarizes the relation between Maynard Smith's position, Sober's 1987 position, and Sober and Wilson's current position.

Which of these three positions is correct? I think Sober and Wilson are right that it is *possible* to apply the concepts of group reproduction and group heritability in the intra-demic selection model. Admittedly these concepts apply more naturally in traditional group selection, where the group-level analogue of organismic reproduction is obvious, and does not need to be established by convention. But there seems no great harm in granting Sober and Wilson their convention for determining group ancestry in the intra-demic model. Nonetheless, I do not think Sober and Wilson's overall position is defensible, nor Maynard Smith's. Strangely enough, I think Sober's 1987 position was actually closest to the mark—that's the position that says that group heritability is *irrelevant* to the issue of group selection. I explain why below.

5 Why the concept of group heritability is irrelevant

Recall Maynard Smith's original intuition: just as individual selection requires heritable variation in the fitness of individual organisms, so group selection requires heritable variation in the fitness of groups of organisms. The dispute so far has focused on the notion of group heritability. But equally crucial is the notion of group *fitness*, which all parties take for granted. In individual selection, an organism's fitness is defined as the expected number of offspring that it will leave. By analogy, the fitness of a group should be taken to mean the expected number of offspring *groups* that it will leave. But that is *not* the concept of group fitness that Sober and Wilson employ. On the contrary, Sober and Wilson equate the fitness of a group with the average *individual* fitness of the organisms in the group ([1998], pp. 26–7, [1994], p. 594). When they talk about the 'differential fitness of groups' as a prerequisite for group selection, they do *not* mean groups which differ in the

number of offspring *groups* they are expected to leave, but rather groups which differ in the average number of *individual* offspring their members are expected to leave.¹³

It is easy to see that these two notions of group fitness—the propensity of a group to leave offspring groups, and the average propensity of individuals in the group to leave individual offspring—are not equivalent. To illustrate, consider two groups A and B each containing 100 organisms. Suppose group A leaves three offspring groups in the next generation, while group B leaves two. By the first criterion of group fitness, group A is fitter. But suppose that the total number of *individual* offspring that members of group A have is 500, and the total number of *individual* offspring that members of group B have is 1000. The average individual fitness of group B is then higher than that of A, and thus B has the higher group fitness, according to the second criterion. So there are two very different notions of group fitness in play.¹⁴

The crucial point is this. Only if group fitness means the propensity of a group to leave offspring *groups*, is the notion of group heritability relevant. If group fitness simply means average individual fitness, then it makes no difference whether there is group reproduction and group heritability or not—what matters is that there should be individual heritability. So Sober and Wilson's mature position contains a serious tension. They go to great lengths to show that the concept of group reproduction and thus group heritability can apply in the intra-demic selection model. But given that the notion of group fitness employed in that model is average individual fitness, the issue of group heritability is simply beside the point.

It is very easy to show that the concept of group heritability is entirely irrelevant to the mechanics of the evolutionary process in which Sober and Wilson are interested. Consider again the example of intra-demic selection from Section 2, where three groups of 100 organisms contribute differentially to the next generation, depending on the proportion of altruists they contain (Table 1). In Section 2, nothing was said about the grouping of the 2nd generation organisms, so no conclusions about group heritability can be immediately drawn.¹⁵ But suppose that the 3600 2nd generation organisms form themselves into nine groups. Suppose further that each of these nine 2nd

¹³ Sober and Wilson actually measure group fitness by the *total* number of individual offspring a group's members have, rather than the average number. But the average number is in fact the relevant quantity in their model. However in their numerical example, Sober and Wilson make the simplifying assumption that all groups are the same size, which means that total number of offspring and average number of offspring can be used interchangeably as a measure of group fitness.

¹⁴ Under certain conditions these two concepts of group fitness may coincide in extension. But the important point is that they need not.

¹⁵ This in itself is actually sufficient to show the irrelevance of group heritability to the intra-demic model, for the process described in Section 2 and summarized in Table 1 is one complete cycle of intra-demic selection.

generation groups contains at least one descendant from each of the three 1st generation groups. Then, by Sober and Wilson’s convention on group reproduction discussed above, *every* first generation group is the parent of *every* second generation group. From this it follows immediately that no group-level trait is heritable. For group heritability means ‘a positive correlation between parental group and offspring groups with respect to a given group-level trait’; and there can obviously be no such correlation, if every 2nd generation group is the offspring of every 1st generation group. Whatever the group-level trait in question, when we plot the value of the trait for parental groups against the average value of the trait for their offspring groups, the result will be a straight line (Figure 3). But group selection has occurred, and will continue to occur, so long as there is positive assortment of altruists in the second generation. So although we can apply the concept of group heritability in the intra-demic selection model if we so choose, whether or not there *is* group heritability makes no difference at all to the workings of that model.

Precisely the same applies to group fitness in the sense of expected number of offspring *groups*. For in the example above, each of the first generation groups has the *same* group fitness in this sense—each leaves nine offspring groups. But none the less, group selection has caused altruism to evolve. Sober and Wilson repeatedly insist that in order for altruism to evolve, we

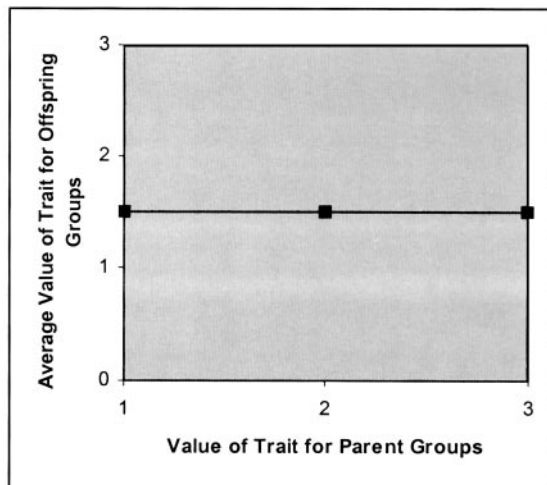


Figure 3: Zero heritability of all group-level traits where every 1st generation group is the parent of every 2nd generation group.

Notes: (i) numerical units are arbitrary; (ii) slope of the correlation function is determined by the fact that every 1st generation group is the parent of every 2nd generation group.

need (a) a set of groups which differ in fitness, and (b) a positive correlation between the fitness of a group and the proportion of altruists it contains. This claim is true, but *only* given that group fitness means average individual fitness. As with group heritability, group fitness in the sense of expected number of offspring *groups* is a parameter that is entirely irrelevant to the process Sober and Wilson are interested in.

To summarize: Sober and Wilson's emphasis on group heritability is beside the point, given that the operative notion of group fitness in the intra-demic selection model is average propensity to leave individual offspring, rather than propensity to leave offspring groups.

This *may* appear to vindicate Maynard Smith's view that intra-demic selection is not real group selection. But in fact it does not. For the notion of group fitness at work in traditional group selection theory was *also* average individual fitness. This crucial point is easy to miss, for two reasons. Firstly, when traditional group selection theorists spoke of adaptations being 'for the good of the group', they *seemed* to mean 'adaptations which benefit the group as a unit', rather than 'adaptations which benefit other individuals in the group'. But when we look at the formal models used in traditional group selection theory, the notion of group fitness actually at work was 'average individual fitness', just as in Wilson's model (e.g. Haldane ([1932]); Levins ([1970])). In such models, the measure of whether an adaptation 'benefits the group' is *not* whether it causes the group to leave more offspring groups, but whether it causes individuals in the group to leave more individual offspring.

Secondly, the role of extinction and colonization in many traditional group selection models (e.g. Levins ([1970]); Gilpin ([1975])), can easily obscure the above point. Colonization and extinction are closely analogous to organismic reproduction and death, fostering the illusion that the groups in group selection play a precisely analogous role to the individuals in individual selection, and thus that group fitness means 'expected number of offspring groups'. But, as Heisler and Damuth correctly observe, differential extinction and colonization are simply processes which contribute to group-level effects on individual fitness ([1987], p. 584). Traditional group selection models which involve extinction and colonization still measure a group's fitness by the number of individual offspring its members leave; this number *may* be enhanced if the group founds many colonies, and it will certainly be reduced if the group goes extinct. But the operative notion of group fitness is *not* expected number of offspring groups. So differential extinction and colonisation are *not* conceptual prerequisites of group selection, *contra* Maynard Smith; they are simply efficient mechanisms for generating groups that differ in their propensity to contribute *individual* progeny to the next generation. Arnold and Fristrup are quite correct that traditional group selection should be 'conceptualised as selection between individuals—but

based on an individual's group membership rather than on characters attributed to that individual' ([1982] p. 299).

It follows that Maynard Smith's position, no less than Sober and Wilson's position, is flawed. To recall, Maynard Smith argued that intra-demic selection should not be co-classified with traditional group selection, because the concept of group heritability applies to the latter but not to the former. This would only be a good argument if traditional group selection theory had defined group fitness as expected number of offspring groups. But as we have seen, it did not. So Maynard Smith's attempt to drive a wedge between intra-demic and traditional group selection fails. Sober's 1987 position is therefore correct: intra-demic selection and traditional group selection belong in the same category, and the issue of group heritability is irrelevant to both.

Obviously, differences between intra-demic and traditional group selection remain—concerning the reproductive isolation, spatial discreteness and duration of the groups. But it would be wrong to see these differences as fundamental. For one thing, reproductive isolation, spatial discreteness and duration are all matters of degree, so any attempt to impose boundaries would invariably be arbitrary. But more importantly, both intra-demic and traditional group selection are concerned with *individual* fitness. Both attempt to model situations where the fitness of an individual cannot be predicted solely on the basis of its own phenotype, but depends on properties of the group to which it belongs. How often such situations arise in nature, and whether they tend to be of the intra-demic or traditional variety, are important empirical questions on which it would be inappropriate to speculate here.

6 On hierarchical application of the principle of natural selection

Advocates of both intra-demic and traditional group selection have often suggested that their theories are the result of applying the abstract principle of natural selection to a higher level of the biological hierarchy than usual—to groups rather than individuals. This is a major theme of Sober and Wilson's book. The foregoing analysis implies that this way of looking at things is most misleading. For by dint of the way they define group fitness, both intra-demic and traditional group selection theorists ensure that there is a fundamental *asymmetry* between the role played by groups in their models and the role played by individuals in models of individual selection. Neither type of group selection constitutes a higher-level analogue of individual selection. This point has occasionally been noticed in the biological literature, but is not widely appreciated. The point is rightly stressed by Heisler and Damuth ([1987]) and Arnold and Fristrup ([1982]).

What would a true higher-level analogue of individual selection look like? Clearly, it would have to posit a population of higher-level entities which 'differ in fitness', and it would have to define the fitness of a higher-level entity as its propensity to leave other higher-level entities as descendants, *not* as the average propensity of its lower-level constituents to leave lower-level descendants. Furthermore, the fitness of a higher-level entity would have to depend on some of its properties, and those properties would have to be heritable, i.e. inherited by its higher-level descendants. Only then would we have a true higher-level application of the principle of natural selection.

Interestingly, the idea of *species selection* defended by macro-evolutionary theorists such as Gould and Eldredge ([1977]), Stanley ([1975], [1979]) and Vrba ([1984], [1989]) satisfies these conditions precisely. As the name suggests, the higher-level entities in species selection are whole species. Species are judged to have 'reproduced' when they give rise to daughter species, i.e. when they speciate. Since different species tend to speciate at different rates, there is differential fitness among species, when species reproduction is understood in this way. Advocates of species selection suggest that the rate at which a given species tends to speciate may depend on properties of the species itself, and that these properties may be passed on to daughter species. Sterelny and Griffiths cite *wide geographical distribution* as a possible species-level property with these characteristics ([1999], p. 206). A widely distributed species is more likely to survive an extinction event, e.g. a meteor striking earth, than a less widely distributed one, and so the probability that it will persist long enough to speciate will be correspondingly higher. Furthermore, it is conceivable that the daughter species to which a widely distributed species gives rise will be widely distributed themselves. If so, then the property of being widely distributed is a species-level adaptation that may in theory have evolved by species selection.

Whether species selection has in fact played an important role in evolutionary history is a highly controversial empirical matter, which I cannot examine here. My point is a purely conceptual one. Unlike group selection, species selection *does* constitute a higher-level analogue of individual selection, for it does *not* define the fitness of a species in terms of the fitness of its component organisms, but rather in terms of the probability of the species founding new species.¹⁶ The dispute between Sober and Wilson and Maynard Smith over group heritability, which I dismissed as irrelevant, would have made perfect sense if the subject under discussion had been species selection. In effect, both Sober and Wilson and Maynard Smith implicitly treat traditional group selection as if it were of the same logical order as species selection, and conduct their discussion of the status of

¹⁶ Arnold and Fristrup are therefore quite correct to say that 'species selection is *not* a higher order form of traditional "group selection"' ([1984], p. 299, authors' emphasis).

intra-demic selection on that basis. But the basis is flawed, for species selection and traditional group selection are of very different logical orders. Sober's ([1987]) claim that group heritability is not necessary for group selection constitutes an implicit recognition of this crucial point.

The contrast with species selection can be used to clarify an issue that I suspect has caused some confusion in the group selection debate. The issue concerns the type of biological adaptation that group selection hypotheses are capable of explaining. Clearly the explanatory target of a hypothesis of species selection is a species-level property—'being widely distributed' in the example above. We postulate a process of between-species selection to explain why current species possess this particular property, which seems to adapt them so well to their environment. In ordinary individual selection, the explanatory target is obviously a property of individual organisms—'the ability to run fast' in the antelope example of Section 1. We postulate a process of between-individual selection to explain why current antelopes possess this particular property, which seems to adapt them so well to their environment. But what about the explanatory target of *group* selection? Are group selection hypotheses supposed to explain properties of groups, or of their component individuals?

The answer to this question is 'of their component individuals'. Both traditional and intra-demic group selection hypotheses explain why it is that current individuals possess a particular property—'the tendency to behave altruistically', in the previous examples. It is very easy to miss this point. For the group selectionist explanation of why there are so many altruists today posits an ancestral population of groups that differed in their proportion of altruists, with a positive correlation between proportion of altruists and group fitness. So it is tempting to think that the explanatory target of this explanation is the group-level property 'proportion of altruists'. This temptation will be especially strong if the *current* population is itself subdivided into groups that differ in their proportion of altruists (which it may or may not be). Anyone who succumbs to this temptation will naturally be led to think of group selection as the higher-level analogue of individual selection. But the temptation is quite misguided: it ignores the fact that the fittest groups in group selection theory are ones which contribute the most offspring *individuals* to the next generation, not the most offspring groups. The explanatory target of group selection hypotheses, whether intra-demic or traditional, is an individual-level property, not a group-level property.

It is possible to imagine a group selection theory that *is* of the same logical order as species selection theory. Such a theory would posit a meta-population of groups differing in their propensity to leave offspring *groups*, where the propensity of any given group to leave offspring groups depends on

a heritable property of that group itself.¹⁷ (In a theory of this type, the existence of group heritability would of course be essential.) Furthermore, the heritable group property in such a theory could even be 'high proportion of altruists'. For it may be that the number of offspring groups a given group leaves is positively correlated with the proportion of altruists in that group; and offspring groups might well resemble their parental groups in respect of proportion of altruists. Such a theory would explain why the current group or groups in the population possess the group-level property 'high proportion of altruists', and would be strictly analogous to species selection. The important point is that neither traditional nor intra-demic group selection is a theory of this type.

In the Introduction, I explained the group selection debate as concerned with the question of whether groups of organisms can be the 'units of selection', but I have not used this expression since then. This was quite deliberate, as 'unit of selection' is not a technical term in evolutionary theory, and has been used in a plethora of different ways by different authors, and by the same authors at different times. Furthermore, the dispute between Sober and Wilson and Maynard Smith does not simply stem from competing stipulations about what 'unit of selection' means, so attempting to resolve the dispute by examining that expression would not have been fruitful. However, my analysis does suggest the following constraint on any coherent definition of 'unit of selection'. If individuals are the units of selection in individual selection theory, and if species are the units of selection in species selection theory, then groups are *not* the units of selection in group selection theory, whether intra-demic or traditional. This may sound paradoxical, but it follows directly from the definition of group fitness in group selection theory as average individual fitness. Formulating a definition of 'unit of selection' that satisfies this constraint is a task I leave for another occasion.

7 Conclusion

Sober and Wilson maintain that the widespread hostility towards group selection in mainstream evolutionary biology is founded on conceptual confusion. The supposed alternatives to group selection, such as kin selection and evolutionary game theory, are actually versions of group selection, they claim. The grounds for this assertion were found to be two: (a) the uncontroversial claim that kin selection and evolutionary game theory are instances of intra-demic selection, and (b) the controversial claim that

¹⁷ In the terminology of Heisler and Damuth ([1987]), a theory of this sort would be 'multi-level selection 2' as opposed to 'multi-level selection 1'. In multi-level selection 1, the fitness of a group is defined in terms of the fitness of its component organisms. Heisler and Damuth correctly observe that virtually all extant group selection theories count as multi-level selection 1.

intra-demic selection *is* group selection. Sober and Wilson's arguments for the latter claim, and Maynard Smith's arguments against, were examined. My analysis shows that Sober and Wilson arrive at a correct conclusion by faulty reasoning. Intra-demic selection and traditional group selection do indeed belong in the same category, but not for the reasons they allege. Their attempt to show that the concept of group heritability can apply in the intra-demic model, and that that model is therefore group selection, is labour in vain. Since intra-demic selection and traditional group selection both define group fitness as average individual fitness, the issue of group heritability is irrelevant to both.

I'll end on a philosophical note. W. V. Quine has argued that empirical and conceptual questions in science are so closely intertwined that it is generally not possible to disentangle them from one another. My analysis of the group selection controversy partially confirms this view. Different authors do use the expression 'group selection' in different ways, and there is no obvious way to classify this disagreement as terminological or substantive, just as Quine says. However, Quine also argues that *because* there is no analytic/synthetic distinction, there is no real room for the traditional philosophical enterprise of trying to clarify scientific concepts, and not much need for it anyway. My analysis disconfirms this view. Subjecting the concepts at work in the group selection debate to philosophical scrutiny is not only possible, but vital if a full understanding of that debate is to be achieved.

Acknowledgements

Thanks to Jonathan Grose, the late W. D. Hamilton, John Maynard Smith, Robert Northcott, Alex Rosenberg, Elliot Sober, David Sloan Wilson and Andy Wells for helpful discussion of these issues, and to two anonymous referees for constructive criticism. This work was carried out while I was the holder of a University of London Jacobsen Research Fellowship, which I gratefully acknowledge.

*Department of Philosophy,
London School of Economics,
Houghton Street,
London WC2A 2AE, UK
s.okasha@lse.ac.uk*

References

- Arnold, A. J. and Fristrup, K. [1982]: 'The Theory of Evolution by Natural Selection: a Hierarchical Expansion', in R. N. Brandon and R. M. Burian (*eds*) [1984], *Genes*,

- Organisms, Populations: Controversies over the Units of Selection*, Cambridge, MA: MIT Press, pp. 292–319.
- Axelrod, R. and Hamilton, W. D. [1981]: 'The Evolution of Co-operation', *Science*, **211**, pp. 1390–96.
- Boorman, S. A. and Levitt, P. R. [1973]: 'Group Selection at the Boundary of a Stable Population', *Theoretical Population Biology*, **4**, pp. 85–128.
- Cassidy, J. [1978]: 'Philosophical Aspects of the Group Selection Controversy', *Philosophy of Science*, **45**, pp. 575–94.
- Charnov, E. L. [1977]: 'An Elementary Treatment of the Genetical Theory of Kin Selection', *Journal of Theoretical Biology*, **66**, pp. 541–50.
- Darwin, C. [1871]: *The Descent of Man and Selection in Relation to Sex*, London: John Murray.
- Gilpin, M. E. [1975]: *Group Selection in Predator-Prey Communities*, Princeton: Princeton University Press.
- Gould, S. J. and Eldredge, N. [1977]: 'Punctuated Equilibrium: the Tempo and Mode of Evolution Reconsidered', *Paleobiology*, **3**, pp. 115–51.
- Grafen, A. [1984]: 'Natural Selection, Kin Selection and Group Selection', in J. Krebs and N. Davies (eds), *Behavioural Ecology: an Evolutionary Approach*, Oxford: Blackwell Scientific Publications, pp. 62–84.
- Haldane, J. B. S. [1932]: *The Causes of Evolution*, London and New York: Harper.
- Hamilton, W. D. [1964a]: 'The Genetical Evolution of Social Behaviour I', *Journal of Theoretical Biology*, **7**, pp. 1–16.
- Hamilton, W. D. [1964b]: 'The Genetical Evolution of Social Behaviour II', *Journal of Theoretical Biology*, **7**, pp. 17–52.
- Heisler, I. L. and Damuth, J. [1987]: 'A Method for Analysing Selection in Hierarchically Structured Populations', *American Naturalist*, **130**, pp. 502–602.
- Levin, B. R. and Kilmer, W. C. [1974]: 'Interdemic Selection and the Evolution of Altruism: a computer simulation study', *Evolution*, **28**, pp. 527–45.
- Levins, R. [1970]: 'Extinction' in M. Gerstenbaher (ed.), 1970, *Some Mathematical Questions in Biology*, Providence, RI: American Mathematical Society, pp. 77–107.
- Lewontin, R. [1970]: 'The Units of Selection', *Annual Review of Ecology and Systematics*, **1**, pp. 1–14.
- Maynard Smith, J. [1964]: 'Group Selection and Kin Selection', *Nature*, **201**, pp. 1145–46.
- Maynard Smith, J. [1976]: 'Group Selection', *Quarterly Review of Biology*, **61**, pp. 277–83.
- Maynard Smith, J. [1982]: *Evolution and the Theory of Games*, New York: Cambridge University Press.
- Maynard Smith, J. [1987a]: 'How to Model Evolution', in J. Dupre (ed.), 1987, *The Latest on the Best: Essays on Evolution and Optimality*, Cambridge, MA: MIT Press, pp. 119–31.
- Maynard Smith, J. [1987b]: 'Reply to Sober', in J. Dupre (ed.), 1987, *The Latest on the Best: Essays on Evolution and Optimality*, Cambridge, MA: MIT Press, pp. 147–50.
- Maynard Smith, J. [1998]: 'The Origin of Altruism', *Nature*, **393**, pp. 639–40.

- Sober, E. [1984]: *The Nature of Selection: Evolutionary Theory in Philosophical Focus*, Cambridge, MA: MIT Press.
- Sober, E. [1987]: ‘Comments on Maynard Smith’s “How to Model Evolution”’, in J. Dupre (ed.), 1987, *The Latest on the Best: Essays on Evolution and Optimality*, Cambridge, MA: MIT Press, pp. 133–45.
- Sober, E. and Wilson, D. S. [1994]: ‘Reintroducing Group Selection to the Human Behavioral Sciences’, *Behavioral and Brain Sciences*, **17**, pp. 585–654.
- Sober, E. and Wilson, D. S. [1998]: *Unto Others: the Evolution and Psychology of Unselfish Behaviour*, Cambridge, MA: Harvard University Press.
- Stanley, S. [1975]: ‘A Theory of Evolution above the Species Level’, *Proceedings of the National Academy of Sciences*, **72**, pp. 646–50.
- Stanley, S. [1979]: *Macroevolution: Pattern and Process*, San Francisco: W. H. Freeman.
- Sterelny, K. [1996]: ‘The Return of the Group’, *Philosophy of Science*, **63**, pp. 562–84.
- Sterelny, K. and Griffiths, P. [1999]: *Sex and Death: An Introduction to Philosophy of Biology*, Chicago: University of Chicago Press.
- Trivers, R. L. [1971]: ‘The Evolution of Reciprocal Altruism’, *Quarterly Review of Biology*, **46**, pp. 35–57.
- Vrba, E. [1984]: ‘What Is Species Selection?’, *Systematic Zoology*, **33**, pp. 318–28.
- Vrba, E. [1989]: ‘Levels of Selection and Sorting with Special Reference to the Species Problem’, *Oxford Surveys in Evolutionary Biology*, **6**, pp. 111–68.
- Wade, M. [1978]: ‘A Critical Review of the Models of Group Selection’, *Quarterly Review of Biology*, **53**, pp. 101–14.
- Williams, G. C. [1966]: *Adaptation and Natural Selection*, Princeton: Princeton University Press.
- Wilson, D. S. [1975]: ‘A Theory of Group Selection’, *Proceedings of the National Academy of Sciences*, **72**, pp. 143–46.
- 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, CA: Benjamin/Cummings.
- Wilson, D. S. [1997]: ‘Introduction: Multilevel Selection Theory Comes of Age’, *American Naturalist*, **150**, supplementary volume, pp. S1–S4.
- Wynne-Edwards, V. C. [1962]: *Animal Dispersion in Relation to Social Behaviour*, Edinburgh: Oliver and Boyd.