Biology and Philosophy **17:** 519–522, 2002.

Reply

Will Peace Follow?

LEE ALAN DUGATKIN
Department of Biology
University of Louisville
Louisville, KY 402092
U.S.A.

Like Kerr and Godfrey-Smith, I am not a pluralist on all issues in evolutionary biology. In many instances one can demonstrate, either theoretically or empirically, that one position on an issue is superior to another in terms of its predictive power. That being said, I simply don't believe this is the case for a comparison of multi-level selection vs. individualist perspectives, and that's why I adopt a pluralist position on this topic. In fact, in 1994, Kern Reeve and I wrote a paper in which we, like Kerr and Godfrey-Smith, showed the equivalence of the multi-level and individualist views, and argued that it is important to be able to translate from one to the other (Dugatkin and Reeve 1994). We claimed that the ability to undertake such translations was important because, in some cases, it will be easier to frame testable predictions using the multi-level perspective, and in other cases it will be easier using the individual perspective.

I applaud Kerr and Godfrey-Smith for extending this view and presenting the argument that in many instances it will be helpful to use both perspectives to understand a single system. In an attempt to adopt Kerr and Godfrey-Smith's approach, I tried to use these two perspectives to understand a system I have worked on for the last fifteen years – predator inspection behavior in fish and the evolution of cooperation.

Pitcher et al. (1986) introduced the term "predator inspection" behavior to describe a fish's slow, jerky movements away from a school and toward a potential predator. This behavior has now been documented in many fish species, and a similar phenomenon exists in both birds and mammals (Dugatkin and Godin 1992a; Dugatkin 1997). On the one hand, a number of studies have shown a cost associated with inspection behavior. For example, inspectors are under greater threat of predation (Dugatkin 1992; Milinski and Parker 1997; but see Godin and Davis 1995) and probably obtain less food than noninspectors (Dugatkin and Godin 1992; Magurran and Girling 1986). Furthermore, Magurran and Higham (1988) have shown that inspectors transfer the information they obtain about a predator to noninspectors (i.e fish that have not incurred the risk associated with inspection). Thus, inspectors pay a cost for inspecting, but the benefits of inspection are shared by all group members.

On the other hand, inspectors do gather important information about putative predators, and the more inspectors present, the less the danger per inspector during sorties toward a predator. In addition, it appears that when no fish inspects, *everyone* in a group faces increased danger (Dugatkin and Godin 1992b). The situation with respect to costs and benefits is even more complex (Godin and Dugatkin, 1997), but for the purposes of this discussion, let's focus on what we have so far.

When thinking about the evolution of inspection (cooperation) in this system, I must admit that using the Class I and Class II visual structures that Kerr and Godfrey-Smith outline for the individualist perspective did indeed help me visualize what the end product in my system should look like. This approach allowed me to easily envision an equilibrium consisting of both inspectors and noninspectors (cheaters). Of course, I could have used a multi-level selection version of Class I and II structures and come up with the same answer. Still, it was easier, at least for me, to do this using the former, rather than the latter, way.

When designing experiments on cooperation and inspection, however, I find the multi-level selection approach most helpful. This approach immediately suggests constructing experiments in which groups differ in the frequency of cooperators, and then subsequent to this, measuring mortality rates both within- and between-groups. When in this mindset, I can readily predict that the equilibrium I envisioned using Class I and II individualist structures will occur as a result of cooperators having lower fitness within groups, but highly cooperative groups having a lower overall mortality rate. In fact, we are indeed running multi-level-selection-inspired trials right now, but the data is still not in.

In the end I would say that I am probably better off having tried to use both the multi-level selection and individualist perspectives in trying to understand inspection and cooperation, and I am certainly no worse for having done so. All that being said, I am not sure exactly what impact the Kerr and Godfrey-Smith paper will have on the study of natural selection. In the first place, while I found this paper interesting, it was rather long and detailed for the point it was trying to get across. What this means is that many people who aren't embroiled in the debate over hierarchical selection probably wouldn't read Kerr and Godfrey-Smith, which will leave only the advocates. Advocates (and I include myself on this list) will have lots to say.

My prediction runs like this: Those who adopt the pluralistic approach presented by Kerr and Godfrey-Smith will have slightly more to work with than they had before, but many will feel that Kerr and Godfrey-Smith wrote an awful lot to get one basic, albeit important, point across. Those who adopt the extreme individualist perspective will again scream bloody murder when they read Kerr and Godfrey-Smith. They will argue that this is the same sort of thing that group selectionists have been saying for years, simply in a slightly different package and with a slightly different twist. Whether or not they will again claim that this is all "a mathematical smokescreen" and a "positively foolhardy" (Dawkins 1979), or that group selection will send evolutionary biology back 40 years (Maynard Smith 1987) is hard to say. But one thing for certain is that they won't be happy. This is a battle they see as over, done with and won, and yet it never seems to just go away.

On the other side of the spectrum, hard-core multi-level selectionists will see Kerr and Godfrey-Smith as yet another sellout, and will argue that Kerr and Godfrey-Smith really are advocating the hierarchical perspective without even knowing it. They'll say that despite the mathematical equivalence of the two perspectives, only multi-level selection explains causation, and that any other way of doing things loses sight of the forest for the trees.

I readily acknowledge that I'm being fairly pessimistic here, and more to the point, that making the sort of predictions I made above is a dangerous activity as the commentaries that run along side mine may show that I am far off base. Perhaps both extremes will view Kerr and Godfrey-Smith's paper as an advance, and those who have long advocated the views presented in the target article will rest easier. If that happens, I will be as happy (and surprised) as Kerr and Godfrey-Smith.

References

Dawkins, R.: 1979, 'Twelve Misunderstandings of Kin Selection', *Z. Tierpsych.* **51**, 184–200.

Dugatkin, L.A.: 1997, *Cooperation Among Animals: An Evolutionary Perspective*, Oxford University Press, New York.

Dugatkin, L.A. and Reeve, H.K.: 1994, 'Behavioral Ecology and "Levels of Selection": Dissolving the Group Selection Controversy', *Advances in The Study of Behaviour* **23**, 101–133.

Dugatkin, L.A. and Godin, J.-G.J.: 1992a, 'Prey Approaching Predators: A Cost-Benefit Perspective', *Ann. Zool. Fennici* **29**, 233–252.

Dugatkin, L.A. & Godin, J.-G.J.: 1992b, 'Predator Inspection, Shoaling and Foraging Under Predation Hazard in the Trinidadian Guppy, *Poecilia reticulata*', *Env. Biol. Fish.* **34**, 265–276.

Godin, J.-G. and Davis, S.A.: 1995, ‘Who Dares, Benefits: Predator Approach Behaviour in the Guppy (*Poecilia reticulata*) Deters Predator Pursuit’, *Proc. R. Soc. Lond.* **259**, 193–200.

Godin, J.-G. and Dugatkin, L.A.: 1997, ‘Female Mating Preference for Bold Males in the Guppy, *Poecilia reticulata*’, *Proc. Nat. Acad. Sci., U.S.A.* **93**, 10262–10267.

Magurran, A.E. and Higham, A.: 1988, ‘Information Transfer Across Fish Shoals Under Predator Threat’, *Ethology* **78**, 153–158.

Maynard Smith, J.: 1987, ‘Reply to Sober’, in J. Dupre (ed.), *The Latest on the Best: Essays on Evolution and Optimality*, MIT Press, Boston, pp. 147–150.

Milinski, M., Luthi, J., Eggler, R. and Parker, G.: 1997, ‘Cooperation Under Predation Risk: Experiments on Costs and Benefits’, *Proc. R. Soc. Lond.* **264**, 831–837.

Pitcher, T.J., Green, D.A. and Magurran, A.E.: 1986, ‘Dicing with Death: Predator Inspection Behavior in Minnow Shoals’, *J. Fish Biol.* **28**, 439–448.

Biology and Philosophy **17:** 523–527, 2002.

Commentary on Kerr and Godfrey-Smith

JOHN MAYNARD SMITH
School of Biogical Sciences
University of Sussex
U.K.

The main thesis of this paper (KGS for short) is that many ‘multilevel’ evolutionary processes can be modelled in two different ways:

- Contextual: fitnesses (expected number of offspring) are assigned to individuals, but depend on what other kinds of individual are present.
- Collective: fitnesses are assigned to groups, but depend on the Proportions of different kinds of individual comprising the group.

They show that, provided clear-cut groups exist, contextual and collective models can both be formulated, and are mathematically equivalent. They do not argue that one method or the other is *in general* superior, but that it may be illuminating to bear both models in mind, both in understanding particular cases, and in seeing analogies between processes.

I have a lot of sympathy with this view although, as explained below, I have some reservations. I agree that equivalent mathematical descriptions are often although not always possible. I also agree that it can sometimes be illuminating to have both models in mind, in two ways. First, a ‘gestalt switch’ between two models can be helpful. For example, the effect of genetic relatedness can be modelled either by ascribing an ‘inclusive fitness’ to individuals (Hamilton 1964), or, as sketched by Hamilton (1963) and made more explicit by Dawkins (1976), by a gene-centred approach: note that, in this example, the individual is the ‘higher level’ entity and the gene the ‘lower level’ one. I confess that in this case I find the gene-centred approach both mathematically simpler and causally more appropriate, but this may merely reflect the fact that I prefer microscopic to holistic models: Maxwell-Boltzmann to classical thermodynamics, and Dawkins to Price’s equation.

But the value of having two types of model is most apparent when comparing different processes. KGS describe the analogy between the stable genetic polymorphism in a single-locus diploid system (the lower level ‘units’ are the alleles and the higher level ‘group’ the individual) and the mixed ESS in the Hawk-Dove game (the individual organisms are the lower level units; it is less clear what is the higher level ‘group’). The recognition of such analogies between levels is particularly helpful when thinking about the 'major transitions' in evolution (Maynard Smith and Szathmary 1995), for example the origin of cells, of eukaryotes, of multicellular organisms, of animal societies. In these transitions, entities which before the transition could reproduce on their own can only do so as part of a larger whole after the transition. The problem is to explain why selection between the lower-level entities for 'selfish' behavior does not disrupt the functioning of the higher level. It turns out that similar mechanisms (relatedness, synergy, policing) are often present at different levels, a fact that only becomes apparent if one sees the formal analogies.

However, I have reservations about the collective approach, which can best be explained in the context of the settlement of animal contests by signalling rather that escalated fighting. I first learnt about this when a zoology undergraduate in the 1940s by reading Konrad Lorennz. I remember that the behaviour was explained by Julian Huxley by saying that ritual evolved because escalated fighting would 'militate against the survival of the species'. Even as a student I knew that had to be wrong. However, I did not think seriously about the problem until 1970, when it was the stimulus for developing evolutionary game theory (Maynard Smith and Price 1973). I saw this as an attempt to replace Huxley's 'group selection' account by an explanation in terms of 'individual selection'. Although animal contest behaviour is still not fully understood, all serious work on the topic is couched in game theoretic terms, in which fitnesses are ascribed to individuals with different behaviours, conditional on what others are doing. I have little sympathy, therefore, with the suggestion that the problem can be treated as an example of group selection. KGS could reply, reasonably, that Huxley's mistake was in his identification of the species as the appropriate 'group'. If the pair of contesting individuals is identified as the group, then a collective model is possible, formally equivalent to the game theoretic model. My difficulty with this is that, although formally possible, such a collective model would be causally inappropriate: I will return to this below.

As always, it is best to start with Darwin. His basic idea can be summarized as follows. Given a population of entities that can reproduce, that have heredity (like begets like), and that vary in ways that affect their chances of reproduction, then the population will evolve characteristics ensuring individual survival and reproduction. This can explain the astonishing adaptedness of living organisms, which otherwise would call for the existence of a divine creator. Unfortunately, by the 1950s it had beconme common for biologists to explain the characteristics of organisms by saying that they contributed to the survival, not of the individual, but of the species, without realizing that such an explanation, although not necessarily wrong, also requires an explanation of why individual selection does not disrupt the species adaptation. Huxley's remark is an example of such thinking: others that stimulated me at the time are Wynne-Edwards' (1962) explanation of social behaviour, and Darlington's (1939) of sex. This facile model of thinking had to be corrected before progress on these topics could be made. Williams (1966), provoked by other examples of 'group' thinking, played a crucial role. I do not share the view of Sober and Wilson (quoted with approval by KGS) that it was 'mistaken' to get rid of this mode of thought: it was essential to do so. But it could not be the end of the discussion, for several reasons. Sex can plausibly be seen as an adaptation enabling species to evolve, and between-species selection has probably been an important cause of its evolution, although certainly not the only one. Selection does sometimes operate between groups of organisms lower than species. Perhaps most important has been a recognition of the distinction between genes as 'replicators' and organisms as 'vehicles' (Dawkins 1982), and the discovery of 'selfish' genetic elements. These problems are today a central concern of evolutionary theorists: this would not be so if 'old-fashioned' group selection thinking, illustrated by Wynne-Edwards and Darlington, had not been banished.

Consider the following list of adaptations in the living world:

i) The halteres of Diptera are sense organs needed to stabilize flight.
ii) The behaviour of worker bees contributes to the productivity of the hive.
iii) The members of a group of meerkats take it in turns to act as sentinels, thus reducing predation on the group.
iv) Signals between funnel-web spiders fighting over webs reduces the likehihood of escalated fights that might injure both parties.
v) In a forest composed of trees mainly of a single species, the morphology of individual trees helps them to compete for light.

How should we think about these examples? Would a pluralist approach be helpful? I do not find it easy to answer these questions, but I will try. The last question is perhaps the easiest. It cannot be treated by the 'collective' approach – or at least not by the one proposed by KGS – because there are no 'groups'. However, the morphology that would maximise the fitness of an individual tree is context-dependent: it depends on how many trees, of what kinds, are its neighbours. I include this example because I think there are many cases in which individual fitness is context-dependent, but there are no definable groups.

It is also easy to answer the first question, but a little harder to justify the answer. We know why flies have halteres: they are an adaptation for flight, and not context-dependent. Most physiological explanations are of this kind: the heart is an 'individual' adaptation, although its optimal size may depend on the amount of conspecific competition. But an individual organism consists of millions of cells, and contains millions of potentially selfish genes. Why do we ignore these facts when thinking about halteres or hearts? Essentially because we treat intercellular and intragenomic conflict as problem in their own right, and assume that they have seen 'solved' when thinking about individual adaptations. I think this is OK, provided we occasionally remember what we are doing.

Questions ii) and iii), concerning bees and meerkats, do seem to me appropriate for a pluralist approach. In both cases, the hard question is why individual organisms behave cooperatively, anyway most of the time. In the case of bees, the answer is complex, involving at least three factors – the sterility of workers, genetic relatedness of workers and queen(s), and 'punishment'. There are many difficulties: for example, workers can lay unfertilized eggs that develop as males, but are 'punished' by other workers if they do so. Obviously, a simple model of the kind proposed by KGS is not adequate, but it is true that there is a 'group' – the colony – whose success contributes to the survival of the genes of the queen that founded it, and that there is room for conflict between the individuals within a hive.

Meerkats also live in clearly defined groups. Relatedness is probably unimportant, and I know of no evidence of the 'punishment' of individuals that do not take their turn as a sentinel. Probably the important factors are that individuals have a better chance of reproducing if they are members of a large group, and that acting as a sentinel may not be as expensive to the individual as might appear at first sight (Clutton-Brock et al. 1999).

I find that the hardest question to answer is the one I have spent most time thinking about – the contest behaviour of spiders, and of animals generally. The contextual approach is obvious, and is the one usually adopted in the form of game theory. Is a collective approach helpful? One can treat a pair of contesting individuals as a 'group', and calculate the 'fitness' of that group as the sum of the fitness, after the contest, of the two individuals, Thus, formally, one could apply the collective model proposed by KGS. But why would one want to?

Although a group of two fighting individuals can be used as an accounting device, it cannot be thought of as a 'cause' of anything. KGS quote Sober and Wilson as saying that 'only a multilevel description is faithful to the causal structure involved.' This seems to me precisely wrong. The cause of what happens is the behaviour of individuals (and, behind that, the genes that cause that behaviour): the fitness of a group is merely the arithmetical sum of the fitnesses of the individuals that compose it. I think that there *are* cases in which a group of individuals can be seen as part of the 'causal structure':

for example, a mated pair of birds both caring for the young may be able to achieve more (e.g. one guards the nest while the other collects food) than could two isolated individuals. But if a pair of individuals are fighting. I cannot see the 'group' of two as constituting a causally effective unit.

The fact that, in some cases, the same process can be described by two mathematically equivalent models, contextual and collective, might suggest that the debates about 'group selection' have been arguments about the use of words rather than about what the world is like. Certainly there has been much semantic confusion: the term 'group selection' has been used to mean everything from the existence of group-level adaptations (e.g. sex) arising by the selective survival of some groups and the extinction of others, to any case in which the fitness of individual organisms is in part determined during interaction with other members of the same species. But I think there are empirical questions as well: in particular, what processes were responsible for the evolution of higher-level entities – cells, organisms and societies?

References

Clutton-Brock. T.H., O'Riain, M.J., Brotherton, P.N.M., Gaynor, D., Kansky, R., Griffin, A.S. and Manser, M.: 1999, 'Selfish Sentinels in Cooperative Mammals', *Science* **284**, 1640–1644.

Darlington, C.D.: 1939, *The Evolution of Genetic Systems*, Cambridge University Press, Cambridge.

Dawkins, R.: 1976, *The Selfish Gene*, Oxford University Press, Oxford.

Dawkins, R.: 1982, *The Extended Phenotype*, Oxford University Press, Oxford.

Hamilton, W.D.: 1963, 'The Evolution of Altruistic Behaviour', *American Naturalist* **97**, 354–356.

Hamilton, W.D.: 1964, 'The Genetical Evolution of Social Behaviour. I, II', *Journal of Theoretical Biology* **7**, 1–16, 17–52.

Maynard Smith, J. and Price, G.R.: 1973, 'The Logic of Animal Colflicts', *Nature* **246**, 15–18.

Maynard Smith, J. and Szathmary, E.: 1995, *The Major Transitions in Evoloution*, Oxford University Press, Oxford.

Williams, G.C.: 1966, *Adaptation and Natural Selection,* Princeton University Press, Princeton.

Wynne-Edward, V.C.: 1962, *Animal Dispersion in Relation to Social Behaviour*, Oliver & Boyd, Edinburgh.

Biology and Philosophy **17:** 529–537, 2002.

Perspectives and Parameterizations

Commentary on Benjamin Kerr and Peter Godfrey-Smith's "Individualist and Multi-Level Perspectives on Selection in Structured Populations"

ELLIOTT SOBER
Philosophy Department
University of Wisconsin
Madison, WI 53706
U.S.A.
E-mail: ersober@facstaff.wisc.edu

DAVID SLOAN WILSON
Biology Department
Binghamton University
Binghamton, NY 13902-6000
U.S.A.
E-mail: dwilson@binghamton.edu

Abstract. We have two main objections to Kerr and Godfrey-Smith's (2002) meticulous analysis. First, they misunderstand the position we took in *Unto Others* – we do not claim that individual-level statements about the evolution of altruism are always unexplanatory and always fail to capture causal relationships. Second, Kerr and Godfrey-Smith characterize the individual and the multi-level perspectives in terms of different sets of parameters. In particular, they do not allow the multi-level perspective to use the individual fitness parameters α_i and β_i. We don't see why the multi-level perspective prevents one from thinking in these terms. Kerr and Godfrey-Smith's argument that Uyenoyama and Feldman's (1980, 1992) definition of altruism belongs more to the individualist perspective than it does to the multi-level perspective is an artifact of their choice of parameters; the same point applies to their argument about the individualism inherent in the idea of Class I and Class II fitness structures.

In *Unto Others* (see, for example, p. 331) we distinguished and endorsed two types of pluralism. First there is a pluralism of perspectives, wherein one can represent a single process in different equally correct ways. Second, there is a pluralism concerning what happens in nature – different types of selection process can and do occur in nature, and they occur in different combinations. Whereas the first type of pluralism brings with it a kind of conventionalism, the second involves a kind of realism. It is a matter of convention which of two equivalent models you use to describe the evolution of a given trait. But it is a factual matter, not a matter of convention, whether the process of group selection occurred in the evolution of that trait, and the same goes for whether individual or genic selection played a role as well. Kerr and Godfrey-Smith (2002) defend a pluralism of the first kind; if this were all they were doing, we would have no disagreement. But, alas, there are some bones that we feel compelled to pick, beginning with the way in which they characterize our position.

At the start of their paper, Kerr and Godrey-Smith describe our position as follows:

> Sober and Wilson agree that it is possible to devise an evolutionary model of altruism that gets the mathematical results right, while only assigning fitnesses to individuals. However, Sober and Wilson insist that only a multi-level description of these systems can yield real *understanding* of the evolution of altruism, because only a multi-level description is faithful to the *causal structure* involved ... The individualist perspective, for Sober and Wilson, is only predictively, and not explanatorily, adequate [italics ours].

The burden of Kerr and Godfrey-Smith's paper is to describe a version of pluralism that they think we reject. They argue that the evolution of altruism can be modeled from the individualist perspective in a way that promotes understanding and correctly describes causal relationships. In fact, we do not disagree with this claim. Their description of our view in terms of "understanding" and "causality" is too vague. We grant that models that are written in what they take to be the language of individualism can be explanatory. To point out, for example, that altruists have higher average fitness than selfish individuals when there is a sufficient correlation among interacting individuals is to make an *explanatory* remark – indeed, it seems natural to say that sufficient correlation is one of the *causes* of the evolution of altruism.

A little later, Kerr and Godfrey-Smith say that they will not address the question "What is group selection and when does it happen?" We want to emphasize that it is this question towards which we take a realist attitude. Indeed, this is the fundamental question we addressed in Part I of *Unto Others*. The basic question at the core of the group selection controversy has always been "can traits evolve by benefiting whole groups, despite being selectively neutral or disadvantageous within groups?" This question leads to the definitions employed by mutli-level selection theory for group selection (which occurs precisely when there is variation in fitness among the groups[1] in a metapopulation) and individual selection (which occurs precisely when there is variation in fitness among individuals within groups). Selection among genes within individuals (aka intragenomic conflict) adds a third level to the hierarchy. If the different types of selection are so defined, it becomes a factual (not a conventional) matter what mix of selection processes influenced the evolution of a given trait.

These definitions are the ones we think are best suited for answering the basic question outlined above. In contrast, there are other semantic frameworks that make it impossible to answer the basic question, however useful they may be in other respects. In particular, if we know only that a trait with higher fitness replaced a trait with lower fitness, we can't tell whether the fitter trait was selectively advantageous within groups, but not between, or was selectively disadvantageous within groups, and evolved on the strength of its being favored by group selection. The information needed for answering the basic question is not present in the statement that the trait that evolved had a higher average fitness than its competitors. Information about these overall averages is like information about net forces; it does not describe the component forces that were at work. The *averaging approach* does not have the resources for answering the basic question about group selection.

Kerr and Godfrey-Smith do not endorse what we have called the *averaging fallacy* (the claim that group selection does not exist *because* fitnesses can be assigned to genes or to traits of individuals by averaging over contexts) and they take pains to point out that their "aim is not to banish or discredit group selection ..." However, we are concerned that their article will lead readers to think that it is a matter of perspective, not of fact, how one answers the basic question "can traits evolve by benefiting whole groups, despite being selectively neutral or disadvantageous within groups?"[2] We hope that Kerr and Godfrey-Smith will make it clear in their reply that the answer to the basic question is "yes" and does not depend on the perspective employed. This is the realism that we are trying to defend.

To summarize our first point – Kerr and Godfrey-Smith think we disagree with the pluralism they endorse. However, they do so based on an incomplete rendering of the kind of realism we favor. Kerr and Godrey-Smith are correct in thinking that we have frequently defended the mutli-level perspective by saying that it promotes understanding by disentangling the different causal processes that lead to an evolutionary outcome (we did so two paragraphs ago). However, we did this in the context of comparing our approach with the *averaging approach* just described. However, we do not claim that *all* individual-level descriptions are unexplanatory and fail to capture causal relationships (see, for example, our remark before about altruism and correlation). Thus, we are realists about the question "when does group selection occur?" We are not realists about the choice between mathematically equivalent models. If two models are equivalent, then either both are true, or neither is, period.

Although we agree with Kerr and Godfrey-Smith that individualism and the multi-level perspective may each be useful, we disagree with how they spell out this distinction. Kerr and Godrey-Smith associate each "perspective" with a particular choice of parameters. As summarized in the first line of their Table 1, they take the individualist perspective to involve using just α_i, β_i, i, and n, and the multi-level perspective to involve using just π_i, ϕ_i, i, and n. We don't see why the multi-level perspective has to be set up in this way. We have often deployed what we think of as a multi-level perspective by talking about selection within groups (in terms of the individual fitnesses α_i, and β_i) and about selection among groups (in terms of the π_is, the group productivities). Kerr and Godfrey-Smith say that the main goal of their paper is to ascertain "... the relation between models that explicitly assign fitness-like parameters to groups, and models that only assign fitness-like parameters to individuals." Of course, a model that assigns *some* parameters to groups need not assign *all* parameters to groups. A multi-level model, properly so called, will assign parameters to objects at more than one level of the biological hierarchy. Our second main point, then, is that we don't see why the multi-level perspective should be associated with the parameterization that Kerr and Godfrey-Smith use.

How, then, should the "multi-level perspective" be construed? As described in Chapter 3 of *Unto Others*, we think of this perspective as providing a heuristic for investigating why different traits evolved. It is often fruitful to ask what form a trait would take if it were shaped purely by group selection, and what form the trait would take if it were shaped purely by individual selection. The empirical question will then be whether the actual state of the trait involves a compromise between these two possible causes. If the trait in question is quantitative, its actual value may fall between these two predictions (perhaps being closer to one than to the other); it also may deviate wildly from both. This heuristic is naturally put to work by thinking about the fitnesses of individuals within the same group (that is, by comparing α_i and β_i, for each i) and also about the fitnesses of groups (represented by the π_is).

Another example of the heuristic value of the multi-level perspective can be found in the example of egg laying in chickens that we described in *Unto Others* (pp. 121–123). Artificial selection for individual hens who laid the most eggs had been the *modus vivendi* for decades, and it had reached a ceiling. It was not until the early 1990s that breeders realized that groups of hens are evolvable units; it was only then that groups of hens were selected for their productivity, with the result that egg yield increased by 160% in six generations. In retrospect, the limitation of the individual selection regime was obvious – in selecting the most productive hens, one was also selecting hens who were able to suppress the egg-laying of their neighbors. Group selection cut through this limitation. The multi-level perspective leads one to think of groups as possible units of selection; breeders caught in the grip of the individualist paradigm were unable to achieve this insight.

Kerr and Godrey-Smith describe two definitions of altruism, and associate one of them with the mutli-level perspective and the other with individualism. The first is the one we used in *Unto Others*.[3] It has two parts:

(8) Altruists are less fit than selfish individuals in the same group.

(9) Groups are fitter the more altruists they contain.

They contrast this with Uyenoyama and Feldman's (1980, 1992) definition, which requires (8) and the following proposition as well:

(10–11) Altruists and selfish individuals in the same group each have higher fitnesses, the more altruists there are in the group.

Kerr and Godrey-Smith are right that neither of these definitions entails the other. For the sake of clarifying their mutual independence, we depict three fitness functions in the accompanying figure. In the lower graph, trait A satisfies the SW definition of altruism, but not the UF definition; in the middle graph, A satisfies the UF definition, but not the SW definition.[4] We've also included a third fitness function in the top graph, in which trait A is altruistic according to both definitions; we'll comment on its significance presently.

In their Table 1, Kerr and Godfrey-Smith point out that the SW definition of altruism takes less ink to write in the mathematical language they associate with the "multi-level perspective" than it does to write in the language they associate with individualism; this point would remain true if the multi-level perspective were allowed to use the individual fitness parameters α_i and β_i. However, Kerr and Godfrey-Smith also note that the UF definition is more concisely expressed in the language they associate with individualism than in the language they associate with the multi-level perspective. This ceases to be true if the multi-level model is allowed to use α_i and β_i; in this case, the two languages express the UF definition in exactly the same way. Kerr and Godfrey-Smith's argument about the individualism of the UF definition of altruism is an artifact of their choice of vocabulary. Whether there is a better argument defending this conclusion remains to be seen; we note here that both definitions describe both individuals and groups, as our verbal rendition of propositions (8), (9), and (10–11) attests.

It is interesting to consider in this connection a definition of altruism that satisfies both the SW and the UF requirements – it endorses all three of propositions (8), (9), and (10–11). Trait A in the top graph in our figure counts as altruistic in this stronger sense. This happens to be the fitness function that is needed to capture what is going on in n-person prisoners' dilemmas, for example. Is this more constrained definition of altruism more a part of individualism or the multi-level point of view? We doubt that it would help answer this question to see how much ink it would take to formulate this idea using the two parameterizations that Kerr and Godfrey-Smith describe. We note, however, that this conjuntion mentions both groups and individuals, just

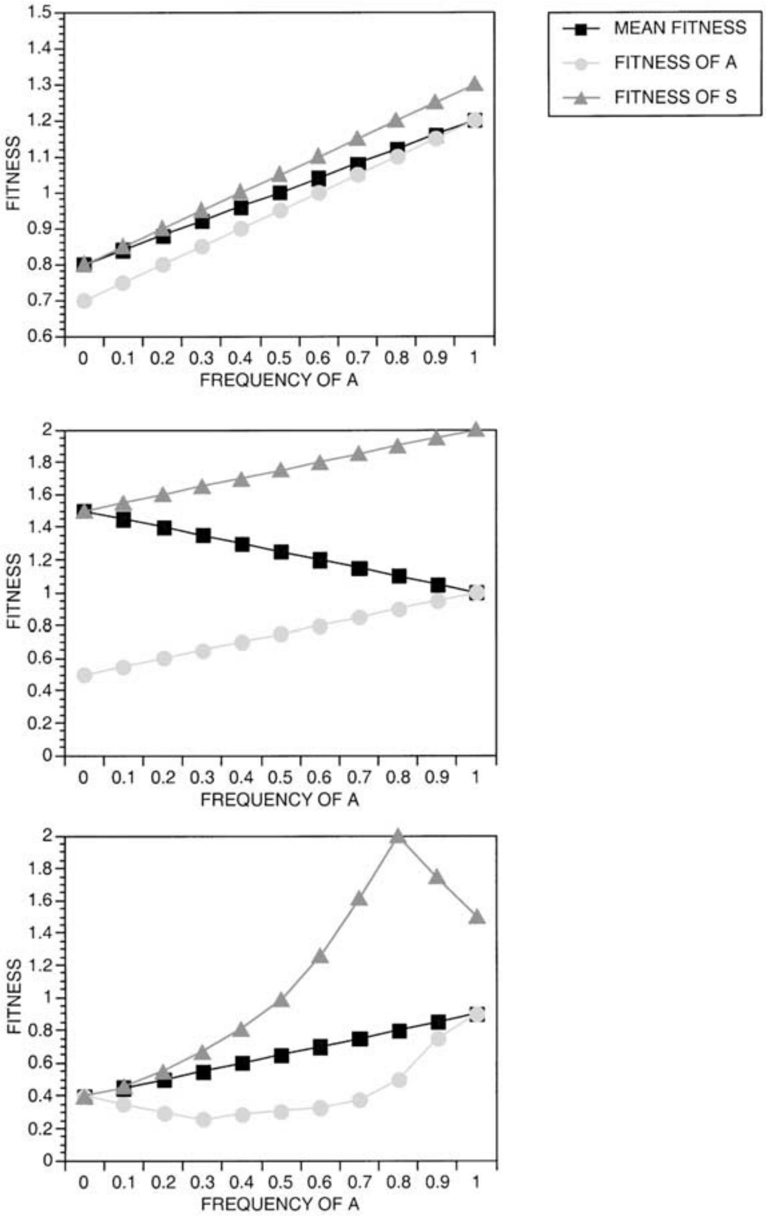

Figure 1. In each of these fitness functions, the fitnesses of traits A and S depend on their frequencies in a group. The $\overline{w}$ line represents the average fitness of individuals in the group, which also is frequency dependent. In the lower graph, trait A counts as altruistic according to the SW definition, but not according to the UF definition; in the middle graph, A is altruistic according to the UF definition, but not according to the SW definition; in the upper graph, A is altruistic according to both definitions.

as the SW and UF definitions do separately. To us, the three fitness functions shown in our figure are equally at home in the multi-level point of view. The intensity of individual selection is depicted by the distance between the S and A lines. The slope of the $\overline{w}$ line is relevant to calculating the intensity of group selection.[5] It is a terminological matter (though one with a history) which of these one uses to define the term "altruism." Since there is more to the question of group selection than the issue of whether altruism evolves, it doesn't much matter to the multi-level point of view which of these definitions one adopts.

Similar remarks apply to the interpretation that Kerr and Godfrey-Smith place on the visual representations given in their Figure 3. They view the first column as part of "the individualist perspective" and the second as part of "the multi-level perspective." We agree that the distinction between lines 1 and 3 is hard to take in at a glance if you use what they call "the multi-level perspective," but is easier to grasp if you use what they call the "individual perspective." However, once again, we don't see why the multi-level perspective can't use the visual representations that Kerr and Godrey-Smith associate with individualism. In saying this, we are not attempting in hindsight to expropriate an alien mode of representation that looks good; both of us have used fitness representations of precisely this sort in many publications, and it is the one we used to construct the figure in the present paper (adding a line to represent $\overline{w}$). Any representation that helps partition within-group from between-group selection is a friend to the multi-level perspective.

We have a similar hesitation about Kerr and Godfrey-Smith's discussion of the distinction between "Class I" and "Class II" fitness structures. These are depicted in their Figure 1. When this distinction is applied to traits that count as altruistic and selfish in the sense of the SW propositions (8) and (9), it coincides with the distinction that one of us (DSW) drew between "strong altruism" and "weak altruism" (Wilson 1990). Weak altruism can evolve when there is random formation of groups, but strong altruism cannot. It will come as no surprise, therefore, that we do not see why the Class I/Class II distinction is any more a part of the individualist perspective than it is part of the multi-level perspective. Kerr and Godfrey-Smith point out in Table 1 that each class takes less ink to define in the mathematical language they associate with the "individualist perspective" than it does to define in the language they associate with the "multi-level perspective." However, if the multi-level perspective is allowed to use α_i and β_i, it can describe each class in precisely the same terms that the individualist perspective uses.

Kerr and Godfrey-Smith use "last member analysis" to characerize the difference between Class I and Class II. They classify this mode of analysis as individualistic; it obtains its intuitive appeal in part by mimicking a reasoning process based on *absolute fitness*: “If I do x, I will receive a fitness of W(x). If I do y, I will receive a fitness of W(y). Therefore I will prefer x to y whenever $W(x) > W(y)$. Hence x will evolve precisely when $W(x) > W(y)$.” This is what one of us called the *heuristic of personification* – one predicts which trait will evolve by thinking about what an individual aiming just to maximize his or her own fitness would choose (Sober 1999). Notice that this reasoning says nothing about *relative* fitness within or among groups and therefore cannot address the basic question that we outlined above. Furthermore, it sometimes yields incorrect predictions about which traits will evolve. Nevertheless, the Matessi-Jayakar proposition is correct in saying that this rule of thumb correctly predicts the outcome of evolution when individuals are randomly distributed into groups. We note, however, that this criterion may have limited usefulness, since so much of the discussion of kin selection and group selection (never mind whether the former is a special case of the latter, as we believe) involves nonrandom group formation; in this case, the criterion simply does not apply. In addition, the MJ result fails to carry over to cases in which individuals are inclined to behave altruistically depending on how many other altruists there are in the group. Consider groups of size N and a range of strategies of the form “be altruistic if fewer than x other individuals in the group are behaving altruistically,” where x can vary from 1 to N. Wilson (1997) showed that the MJ proposition and last member analysis fails to predict the outcome of evolution in randomly formed groups. The reason is that an individual who decides to become an altruist does not necessarily add an altruist to the group, because another individual might decide to turn selfish based on its strategy and the focal individual’s decision. Stated in multi-level terms, in randomly formed groups, the interactions among the strategies result in below-random phenotypic variation among groups, decreasing the importance of among-group selection. More generally, rules of thumb based on absolute fitness frequently do not substitute for an explicit multi-level selection model, even with random group formation.

In summary, our main reservation about Kerr and Godfrey-Smith’s approach is that it depends so much on defining the individualist and multi-level perspectives in terms of two specific parameterizations. In particular, we don’t see why the multi-level perspective can’t use the individual fitness parameters α_i and β_i. However, if the multi-level perspective gets to use all the language that the individualist perspective gets to use, and more, doesn’t this show that the individualist perspective has no heuristic value? We don’t think so. For example, we regard the stability condition involved in the idea of an *evolutionary stable strategy* as individualist in character; it says that trait A is an ESS (where B is the alternative trait) if and only if $w(A, A) > w(B, A)$.

An A individual has to do better when paired with another A individual than a B individual does when paired with an A. It doesn't matter how much ink this idea takes to write down in different languages. The relevant fact is that it involves thinking about a single organism in one context and seeing how its fitness compares to another single organism in a different context. Sometimes thinking just about organisms is heuristic, but at other times thinking about organisms and groups separately is heuristic.

Acknowledgments

We thank Benjamin Kerr and Peter Godfrey-Smith for useful discussion.

Notes

[1] Of course, this definition of group selection requires elucidation of what counts as a "group," on which see *Unto Others*, pp. 92–98.

[2] There is a trivial sense in which every interrogatory sentence depends for its answer on a convention – on the conventional assignment of meanings to words. Since this is a universal feature of language, it has nothing in particular to do with the units of selection problem.

[3] The numbered propositions that follow are English translations of propositions with the same numbering in Kerr and Godfrey-Smith's paper.

[4] Notice that the A trait cannot evolve if it obeys fitness function in our middle graph, since the trait will be selected against both within groups and among groups.

[5] More information is needed, however, concerning how groups are formed, if one is to determine the intensity of group selection.

References

Kerr, B. and Godfrey-Smith, P.: 2002, 'Individualist and Multi-Level Perspectives on Selection in Structured Populations', *Biology and Philosophy* **17**, 477–517.

Sober, E.: 1999, 'Three Differences Between Evolution and Deliberation', in P. Danielson (ed.), *Modeling Rationality, Morality, and Evolution*, Oxford University Press, New York, pp. 408–422.

Sober, E. and Wilson, D.S.: 1998, *Unto Others – the Evolution and Psychology of Unselfish Behavior*, Harvard University Press, Cambridge.

Uyenoyama, M. and Feldman, M.W.: 1980, 'Theories of Kin and Group Selection: A Population Genetics Perspective', *Theoretical Population Biology* **17**(5), 380–414.

Wilson, D.S.: 1990, 'Weak Altruism, Strong Group Selection', *Oikos* **59**, 135–140.

Wilson, D.S.: 1998, 'Hunting, Sharing and Multilevel Selection: The Tolerated Theft Model Revisited', *Current Anthropology* **39**, 73–97.