Multilevel Selection and the Major Transitions in Evolution

Samir Okasha†‡

A number of recent biologists have used multilevel selection theory to help explain the major transitions in evolution. I argue that in doing so, they have shifted from a 'synchronic' to a 'diachronic' formulation of the levels of selection question. The implications of this shift in perspective are explored, in relation to an ambiguity in the meaning of multilevel selection. Though the ambiguity is well-known, it has never before been discussed in the context of the major transitions.

1. Introduction. The debate in evolutionary biology over 'levels of selection', long of interest to philosophers of science, has undergone a subtle transformation in recent years. The transformation coincides, more or less, with the increasing prominence of 'multilevel selection theory' in the evolutionary literature. The basic idea of this theory—that natural selection may operate simultaneously at more than one level of the biological hierarchy—is not in itself new; indeed, it can be traced back to Darwin himself, and forms the basis for the traditional levels of selection discussions of the 1960s and 1970s. What is new is the *use* to which multilevel selection is currently being put. Increasingly, biologists interested in explaining what Maynard-Smith and Szathmáry call the 'major transitions in evolution' have made use of ideas from multilevel selection theory (Buss 1987; Michod 1997, 1999; Maynard Smith and Szathmáry 1995; Michod and Roze 1999; Roze and Michod 2001; Frank 1998). The work of these authors extends the traditional levels of selection question in an important new way.

The 'major transitions in evolution' refer to the transitions from solitary

†To contact the author, please write to: Department of Philosophy, University of Bristol, 9 Woodland Road, Bristol, BS8 1TB United Kingdom; e-mail: samir.okasha@bris.ac.uk.

‡Thanks to Peter Godfrey-Smith, Kim Sterelny, Brett Calcott, Susanna Rinard, Rick Michod, Alirio Rosales, John Damuth, and Jim Griesemer for comments and discussion.

Philosophy of Science, 72 (December 2005) pp. 1013–1025. 0031-8248/2005/7205-0030$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/508102 Published online by Cambridge University Press

replicators to networks of replicators enclosed within compartments, from independent genes to chromosomes, from prokaryotic cells to eukaryotic cells containing organelles, from unicellular to multicellular organisms, and from solitary organisms to colonies. Some of these transitions occurred in the distant evolutionary past, others much more recently. In each case, a number of smaller units, originally capable of surviving and reproducing on their own, became aggregated into a single larger unit, thus generating a new level of biological organization. The challenge is to understand these transitions in Darwinian terms. Why was it advantageous for the lower-level units to sacrifice their individuality, cooperate with one another, and form themselves into a larger corporate body? And how could such an arrangement, once first evolved, be evolutionarily stable?

This is where multilevel selection theory enters the picture. As Buss, Michod, and Maynard Smith and Szathmáry all stress, to understand the major transitions we need to know why lower-level selection did not disrupt the formation of the higher-level unit. In the transition to multicellularity, for example, we need to know why selection between competing cell lineages did not disrupt the integrity of the emerging multicellular organism. One possibility is that selection acted on the higher-level units themselves, leading them to evolve adaptations that minimize conflict and increase cooperation among their constituent parts. Thus in the case of multicellularity, Buss and Michod argue that early sequestration of the germ-line may be one such adaptation, for it reduces the probability that mutant cells, arising during ontogeny, will find their way into the next generation. Another idea is that passing the life-cycle through a single-celled stage, as occurs in most animal and plant species, is an adaptation for minimizing within-organism conflict, for it increases the relatedness, hence decreases the competition, between the cells within an organism. These particular examples have both been contested, but the general idea that the major transitions involve an interaction between selection at different levels is very widely accepted.

2. Synchronic versus Diachronic Formulations of the Levels Question. Though the theoretical developments reviewed above are still in their infancy, one important conceptual moral is already clear: the traditional formulations of the levels of selection question, though not wrong, are insufficiently *general*, given the use to which multilevel selection is now being put. Traditionally the question has been formulated roughly as follows: the biological world is hierarchically organized—genes are nested within chromosomes, chromosomes within cells, cells within organisms, organisms within groups etc. Furthermore, natural selection will operate on any entities that exhibit 'heritable variation in fitness'. Since entities

https://doi.org/10.1086/508102 Published online by Cambridge University Press

at many different hierarchical levels satisfy this condition, there is the potential for selection to operate at multiple levels.

The problem with formulations of this type, as Griesemer (2000) insightfully notes, is that it takes the existence of the biological hierarchy for granted, as if hierarchical organization were simply an exogenously given fact about the organic world. But of course the biological hierarchy is *itself* the product of evolution—entities further up the hierarchy, such as multicellular organisms, have obviously not been there since the beginning of life on earth. The same is true of cells and chromosomes. So ideally, we would like an evolutionary theory that explains how the biological hierarchy came into existence, rather than treating it as a given. From this perspective, the levels of selection question is not simply about identifying the hierarchical level(s) at which selection *now* acts, which is how it was traditionally conceived, but about identifying the mechanisms that led the hierarchy to evolve in the first place.

This new 'diachronic' perspective gives the levels of selection question a renewed sense of urgency. Some biologists were inclined to dismiss the traditional debate as a storm in a teacup—arguing that in practice, selection on individual organisms is the only important selective force in evolution, whatever about other theoretical possibilities. But as Michod (1999) stresses, multicellular organisms did not come from nowhere, and a complete evolutionary theory must surely try to explain how they evolved, rather than simply taking their existence for granted. So levels of selection other than that of the individual organism *must* have existed in the past, whether or not they still operate today. From this expanded point of view, the argument that individual selection is 'all that matters in practice' is clearly unsustainable.

The contrast I have drawn between the modern 'diachronic' view of the levels of selection question and the traditional 'synchronic' view should not be overdone. Even in the earlier discussions, it was recognized that entities at different levels of hierarchical organization form a temporal sequence, i.e., lower-level entities generally evolved before higher-level ones. But the relevance of multilevel selection to the major transitions was not widely appreciated until Buss's seminal book, *The Evolution of Individuality*. Consider for example Richard Dawkins' brief discussion of how independently replicating units may originally have become grouped into compartments. Dawkins says that it is 'easily understood' why independent replicators might have gained an advantage by 'ganging up together' into cell-like compartments, because their biochemical effects might have complemented each other (1982, 252). This is a plausible idea, but what Dawkins fails to realize is that it in effect invokes group selection. From the selective point of view, replicating molecules combining themselves into compartments is strictly analogous to individual organisms

https://doi.org/10.1086/508102 Published online by Cambridge University Press

combining themselves into colonies or groups. But Dawkins is an implacable *opponent* of group selection, insisting on the impotence of selection for group advantage as an evolutionary mechanism! Clearly, Dawkins has failed to realize that trying to explain the major transitions involves us in levels of selection issues closely analogous to those involved in the traditional debate over group selection.

Another respect in which the traditional discussions sit uneasily with the new diachronic perspective concerns the very *concepts* used to understand selection at multiple levels. Consider for example the famous 'replicator-interactor' conceptualization of natural selection, popularized by Dawkins and Hull, and widely used in the early levels of selection debates. Replicators are defined as entities that 'pass on their structure intact' and are characterized by their 'longevity' and 'copying fidelity'; interactors are entities that 'interact as cohesive wholes with their environment' in a way that causes replication to be differential (Dawkins 1976; Hull 1981). This account of natural selection proved valuable for certain purposes, but it is not well-suited to studying the major transitions, for a simple reason. The longevity and copying fidelity of replicators (such as genes) and the cohesiveness of interactors (such as organisms) are highly *evolved* properties, themselves the product of many rounds of cumulative selection (Griesemer 2000). The very first replicators must have had extremely *poor* copying fidelity (Maynard Smith and Szathmáry 1995), and the earliest multicellular organisms must have been highly *non*cohesive entities, owing to the competition between their constituent cell-lineages (Buss 1987; Michod 1999). If we wish to understand how copying fidelity and cohesiveness evolved in the first place, we clearly cannot build these notions into the very concepts used to describe evolution by natural selection.

These reflections suggest that the diachronic formulation of the levels of selection question ('how did the levels in the biological hierarchy evolve initially?') differs substantially from the synchronic formulation ('at which hierarchical levels does selection now act?'). If this is correct, it raises an overarching question. What becomes of the issues traditionally discussed under the 'levels of selection' rubric, when we move from a synchronic to a diachronic perspective? In what follows, I examine one such issue: the well-known ambiguity in the concept of multilevel selection discussed by Arnold and Fristrup (1982), Sober (1984), Damuth and Heisler (1988), Okasha (2001, 2004), and others. I shall argue that this issue looks substantially different when considered in relation to the major transitions.

3. The Two Types of Multilevel Selection. The basic principle of natural selection is extremely simple: if variant entities in a population experience differential reproductive success, and if the variation is heritable, then the

https://doi.org/10.1086/508102 Published online by Cambridge University Press

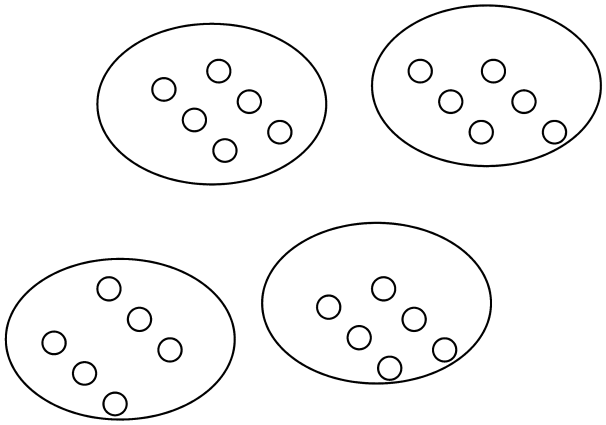


Figure 1. Particles nested within collectives.

composition of the population will change over time. Somewhat surprisingly, given the simplicity of this basic principle, the concept of multilevel selection contains a crucial ambiguity. The ambiguity has usually been discussed in relation to individual and group selection, but it generalizes to any multilevel scenario. Damuth and Heisler (1988) offer a particularly clear discussion of the two meanings of multilevel selection, so I follow their treatment here.

Consider the abstract multilevel scenario depicted in Figure 1, where a number of lower-level units ('particles') are nested within higher-level units ('collectives') for at least part of their life-cycle. The key issue is whether the particles or the collectives constitute the 'focal' level, i.e., the level of interest. Are we interested in the frequency of different types of particle in the overall population of particles (which just so happens to be subdivided into collectives), or are we interested in the frequency of different types of collective themselves? If the former, then the collectives are significant only in that they constitute an environment for the particles, and thus may affect the particles' fitnesses. We will judge evolution to have occurred when the overall frequency of different particle-types has changed. If the latter, then we are interested in the collectives in their own right, not simply as environments of the particles. We will judge evolution to have occurred when the frequency of different collective-types has changed. Damuth and Heisler (1988) refer to the first approach as multilevel selection 1 (MLS1), and the latter as multilevel selection 2 (MLS2).

From the formal point of view, the difference between MLS1 and MLS2 concerns the relation between the fitnesses at the two levels. In both MLS1

https://doi.org/10.1086/508102 Published online by Cambridge University Press

and MLS2, particle fitness has the same meaning: a particle's fitness is the expected number of offspring particles it leaves. Not so with collective fitness. In MLS1, a collective's fitness is defined as the average fitness of the particles within the collective. So the fittest collectives, in this sense, are those which contribute the most *particles* to the next generation (per capita). In MLS2, a collective's fitness is defined differently, as the expected number of offspring *collectives* contributed to the next generation. So the fittest collectives, in this sense, are those that contribute the most offspring *collectives* to the next generation, irrespective of how many particles those offspring collectives contain. In short, in MLS1 we are interested in counting particles while in MLS2 we are interested in counting collectives; thus the meaning of 'collective fitness' is quite different in each.

Most of the literature on 'group selection' has been concerned with MLS1: the aim has been to understand the evolution of an *individual* phenotype, often altruism, in a population subdivided into groups. Hence in formal models of group selection, 'group fitness' has generally been defined as the average fitness of the individuals in the group. By contrast, the macroevolutionary literature on species selection has been concerned with MLS2: the aim has been to understand the changing frequency of different types of *species*, not individuals, over geological time. Hence 'species fitness' is *not* defined as average individual fitness, but rather as expected number of offspring species. The fittest species are not the ones whose constituent organisms are especially well-adapted, but rather the ones with the greatest probability of surviving and speciating (i.e., reproducing). Damuth and Heisler argue that a failure to distinguish clearly between MLS1 and MLS2 plagued many traditional discussions of the levels of selection; see Okasha (2001) for extended discussion of this point.

There is in fact a general reason why the distinction between MLS1 and MLS2 has not always been heeded. For in many circumstances the existence of one type of multilevel selection will imply the other. The collective that contributes the most particles to the next generation, and so is fittest by the MLS1 criterion, will often contribute the most collectives to the next generation, and hence be fittest by the MLS2 criterion too. (This will always be so if the number of particles per collective is constant.) So ordering the collectives by their MLS1 fitness and by their MLS2 fitness will often give the same result. However, this does not undermine the importance of the distinction, for two reasons. Firstly, there is no *necessity* that either type of multilevel selection will imply the other—it depends on contingent biological facts. Secondly, even in cases where MLS1 and MLS2 do go hand-in-hand, the *absolute* fitness of a collective in the MLS1 sense is different from its absolute fitness in the MLS2 sense, for these quantities are measured in different units (number of offspring particles and number of offspring collectives, respectively). So while the

https://doi.org/10.1086/508102 Published online by Cambridge University Press

relative fitnesses of collectives in MLS1 and MLS2 may be identical—relative fitness being a dimensionless quantity—the conceptual difference between the two types of multilevel selection remains. For the absolute fitnesses will never be identical, and absolute fitness is the conceptually prior notion.

One final point of difference between the two types of multilevel selection is this. In MLS2 it is essential that the collectives themselves form parent-offspring lineages. In MLS1 on the other hand, only the lower-level particles need form parent-offspring lineages. This follows from the different meaning of collective fitness in each (Okasha 2001). Thus in MLS2, the collectives must be sufficiently cohesive and discrete for us to discern patterns of ancestry and descent among them. For any given collective, it must be possible to decide which parental collective(s) it came from, and which offspring collectives it produced; otherwise, collective fitness in the MLS2 sense cannot meaningfully apply. In MLS1 on the other hand, where collectives are simply part of the particles' environment, there is no such requirement.

The MLS1/MLS2 distinction was formulated with a synchronic version of the levels of selection question in mind. But what becomes of the distinction from a diachronic perspective? If we wish to use multilevel selection theory to explain the major transitions in evolution, as many biologists believe we should, is MLS1 or MLS2 better suited to the job?

A natural first response is to say that MLS2 is the type of multilevel selection relevant to the major transitions. For MLS1 can only explain particle-level adaptations (or the change in frequency of particle-level traits), not collective-level adaptations. And in a major transition, we are interested in understanding how new collectives, e.g., multicellular organisms, came into existence, from an ancestral state in which they did not exist. On the standard view, this required the collectives to evolve adaptations for reducing conflict between their constituent particles (e.g., germ-line sequestration). These adaptations are properties of the collectives themselves, not the particles. So we are interested in the collectives in their own right, not simply as environments of the particles. Hence collective-level selection must mean selection between collectives based on differential production of offspring *collectives*, not offspring particles.

This argument *sounds* right, and in part it is. But it is not the whole story, for it only applies to the *later* stages of a major transition, when the collectives have already evolved as discrete units, capable of having fitnesses in the MLS2 sense. In the *early* stages of a transition, matters are different. These early stages are characterized by the spread of cooperation among the particles, a prelude to the day when they will sacrifice their individuality entirely and form cohesive collectives. So we should expect traditional theories for the evolution of cooperation/altruism to be

https://doi.org/10.1086/508102 Published online by Cambridge University Press

applicable to the early stages of a major transition. One such theory, of course, is group selection. And as noted above, most group selection models for the evolution of cooperation/altruism are of the MLS1 type: they are concerned with the spread of an *individual* phenotype in a group-structured population, and hence define group fitness as average individual fitness. So although in a major transition we are certainly interested in the collectives in their own right, and not simply as environments of the particles, the *early* stages of a transition must be characterized by the spread of cooperation among the particles, and the type of multilevel selection relevant to this is MLS1.

If this is right, it suggests that MLS1 and MLS2 may both be relevant to the major transitions, each at a different stage. An examination of the biological literature confirms this argument.

4. The Major Transitions and the MLS1/MLS2 Distinction. Michod (1999), Szathmáry and Demeter (1987) and Maynard Smith and Szathmáry (1995) have all discussed one of the very earliest transitions—from independent replicators to networks of replicators encapsulated in a protocell. The first stage in this transition, these authors argue, was the evolution of RNA molecules that code for replicase. This is an altruistic trait, for by coding for replicase an RNA molecule reduces its own replication speed but increases that of other RNA molecules. So selfish RNAs that do not code for replicase should spread faster. How then can replicase coding evolve? This is clearly analogous to the problem of how altruistic behavior can evolve among animals. Michod (1999) suggests that Wilson's (1975) 'trait group' model may provide the key; Maynard Smith and Szathmáry (1995) explore a similar idea. If RNA molecules on a surface become clustered together into groups, then groups containing lots of replicase coding molecules may out-reproduce groups containing fewer, even though *within* each group, selfish nonreplicase coding molecules have an advantage. In this way, group selection may lead the replicase coding RNA molecules to evolve.

There are two crucial points to note here. Firstly, the 'groups' of RNA molecules, that initially consist of independent molecules interacting with each other, are the precursors of the compartmentalized proto-cells that eventually evolve. Secondly, the type of group selection involved is MLS1: group fitness is defined as average individual fitness. Groups containing lots of altruists 'out-reproduce' other groups in that they contribute more *individuals* to the next generation. So what the trait group model explains is the spread of an individual trait—altruism—in the overall population, not the spread of a group trait. This is true both of Wilson's original formulation of the trait group model and of its applications to prebiotic evolution by Michod, Maynard Smith, Szathmáry and Demeter. This

https://doi.org/10.1086/508102 Published online by Cambridge University Press

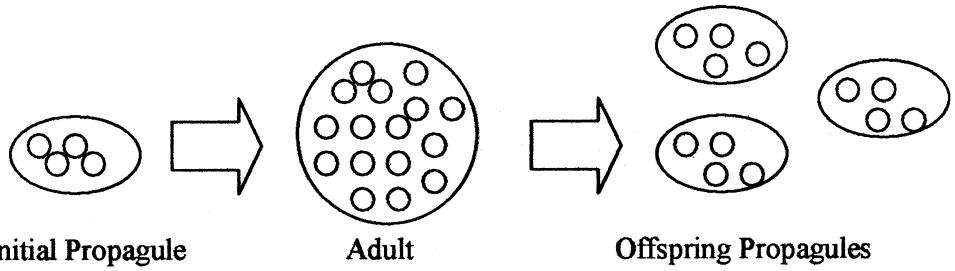


Figure 2. Michod's model for the evolution of multicellularity; $N = 4$.

demonstrates that, at least in the early stages of a major transition when the higher-level collective is still in the process of emerging, multilevel selection of the MLS1 type is relevant. For although MLS1 can only explain facts about the frequency of particle traits, not collective traits, the earliest stages of a major transition are precisely characterized by the spread of a particle trait—cooperation. For groups of cooperating particles represent the first stage in the transition to a new collective.

Michod's more recent work on the evolution of multicellularity illustrates how multilevel selection of the MLS2 variety is also relevant to the major transitions (Michod 1997, 1999). This work is based around a simple multilevel model with the following life-cycle. A multicellular organism (or proto-organism) begins life as a propagule containing N cells, all of which come from the same parent. (If $N = 1$, then the life-cycle starts from a single-celled stage; if $N > 1$ then reproduction is vegetative.) The propagule develops and grows by cell division, reaches adulthood, and then gives rise to new propagules of its own (see Figure 2). There are two levels of selection in this model: between cells within organisms (owing to different cell-types dividing at different rates) and between organisms (owing to different organism-types producing different number of propagules). The overall evolutionary dynamics of the system depends on the balance between the two selective forces. This is an MLS2 model, for the following reason: organism fitness and cell fitness are defined in different units, and across different time scales. A cell's fitness is defined as its rate of division within the organism's life-span, while an organism's fitness is defined as its rate of propagule production. Hence an organism's fitness is *not* equal to the mean fitness of the cells within it, although these two quantities may on some occasions be proportional to each other. This is the defining mark of an MLS2 theory.

Michod and Nedelcu (2003) make some interesting remarks about the relation between the fitnesses at the two levels, in this model. They argue

https://doi.org/10.1086/508102 Published online by Cambridge University Press

that during a major transition, the fitness of the higher-level unit must be 'decoupled' from the fitness of the lower-level units of which it is composed: "group fitness, is, initially, taken to be the average of the lower-level individual fitnesses; but as the evolutionary transition proceeds, group fitness becomes decoupled from the fitness of its lower-level components" (Michod and Nedelcu 2003, 66). This remark bears directly on the MLS1/MLS2 distinction. For in effect, Michod and Nedelcu are saying that in the early stages of a transition, collective fitness is defined in the MLS1 way, but as the transition proceeds, collective fitness becomes 'decoupled' from particle fitness and defined in the MLS2 way.

What exactly do Michod and Nedelcu mean by 'decoupling' collective fitness from particle fitness? This is actually less simple than appears at first sight. For they do not *just* mean defining collective fitness in the MLS2 rather than the MLS1 way, though this is part of it. 'Decoupling' means not just that collective fitness not be *defined* as average particle fitness, but also that collective fitness not be *proportional* to average particle fitness. Michod and Roze's (1999) discussion of the multicellularity model illustrates this point. As noted above, in this model organism fitness is defined as expected number of offspring propagules, *not* as average cell fitness, so qualifies as MLS2. However, Michod and Roze note that in the early stages of multicellularity, all differences in organism fitness may stem solely from differences in adult size—since bigger adults can leave more propagules. Moreoever, differences in adult size will themselves stem from differences in cellular composition—since fast-dividing cells make bigger adults. So in these early stages, an organism's fitness will be strictly proportional to the average fitness of the cells in the propagule from which the organism came. An organism whose constituent cells are very fit, i.e., divide very fast, will grow to a bigger size, hence leave more offspring propagules. So differences in organism fitness are determined by differences in average cell fitness. Proper 'decoupling' only occurs, Michod and Roze argue, when organismic fitness ceases to be dependent (or at least solely dependent) on adult size, hence ceases to be dependent on the division rate of the cells in the organism. Only once this has been achieved has fitness truly 'emerged' at the higher level.

What this suggests is that the relation between particle fitness and collective fitness *itself* undergoes a transformation, during the major transitions. The very earliest stages of a major transition are characterized by interactions among the lower-level particles. Cooperative particles must spread at the expense of selfish ones, which requires that cooperative particles must be clustered together in groups. Multilevel selection of the MLS1 sort is applicable here. As the major transition proceeds, collective fitness must be 'decoupled' from particle fitness. This itself occurs in two stages. Firstly, collective fitness must be defined in the MLS2 way, i.e.,

https://doi.org/10.1086/508102 Published online by Cambridge University Press

we must start to count the collectives themselves, not just their constituent particles. This requires that the collectives be sufficiently discrete to form identifiable ancestor-descendent lineages. Secondly, collective fitness must become independent of particle fitness, in that the fitness of a collective must cease to be proportional to the mean fitness of the lower-level particles. (For otherwise differences in collective fitness are mere side effects of differences in particle fitness.) Once this final stage is completed, then full decoupling of collective fitness from particle fitness has been achieved, and MLS2 occurs autonomously of MLS1. This three stage process is shown below.

Relation between Collective Fitness and Particle Fitness during the Major Transitions.

Stage 1. Collective fitness defined as average particle fitness (cooperation spreads among particles).

Stage 2. Collective fitness not *defined* as average particle fitness, but still *proportional* to average particle fitness (collectives start to emerge as entities in their own right).

Stage 3. Collective fitness neither defined as nor proportional to average particle fitness (collectives have fully emerged; fitnesses are decoupled).

Stage 2 of this process, where the fitness of the higher-level collective is not defined as but is proportional to mean particle fitness, is interestingly reminiscent of an issue familiar from the species selection literature of the 1980s. Vrba (1984) and Vrba and Gould (1986) argued that many alleged cases of species selection in fact involved only 'species sorting'. The latter occurs when species experience differential rates of extinction as a result of differences in the fitnesses of their component organisms. In such cases all the selection is at the organismic level, according to Vrba: the differences in species fitness are mere side effects, or epiphenomena, of lower-level processes. In species sorting as characterized by Vrba, the relation between the higher and lower-level fitnesses is exactly as in Stage 2 above: species fitness is not *defined* as mean organismic fitness, but it is proportional to it. If Vrba is right that species sorting involves only lower-level selection, parity of argument would suggest that the same applies to Stage 2 of a major transition. Interestingly, Michod and Roze (1999) toy with precisely this suggestion in relation to the evolution of multicellularity model. But whether Vrba *is* right is a further question that cannot be resolved here.

5. Conclusion. If the foregoing arguments are correct, then both types of

https://doi.org/10.1086/508102 Published online by Cambridge University Press

multilevel selection are relevant to the major transitions. MLS1 is relevant to the early transitional stages, where cooperation spreads among the particles. This is essential if the collectives are to achieve the degree of cohesion necessary for collective fitness in the MLS2 sense to be a meaningful notion. The key point is that the relation between lower and higher-level fitnesses *itself* undergoes a change, during a major transition.

I suspect that there is a general moral here. The conceptual issues that form the core of the traditional levels of selection debate, of which the MLS1/MLS2 ambiguity is but one, are subtly transformed when we move from a synchronic to a diachronic formulation of the levels question, as we must do if we are to understand the major transitions in Darwinian terms. But the transformation is not so drastic that the traditional discussions lose all their relevance.

REFERENCES

Arnold, Anthony J., and Kurt M. Fristrup (1982), "The Theory of Evolution by Natural Selection: A Hierarchical Expansion", *Paleobiology* 8: 113–129.

Buss, Leo (1987), *The Evolution of Individuality*. Princeton, NJ: Princeton University Press.

Damuth, John, and Lorraine Heisler (1988), "Alternative Formulations of Multilevel Selection', *Biology and Philosophy* 3: 407–430.

Dawkins, Richard (1976), *The Selfish Gene*. Oxford: Oxford University Press.

——— (1982), *The Extended Phenotype*. Oxford: Oxford University Press.

Frank, Steven (1998), *Foundations of Social Evolution*. Princeton, NJ: Princeton University Press.

Griesemer, James (2000), "The Units of Evolutionary Transition", *Selection* 1: 67–80.

Hull, David (1981), "Units of Evolution: A Metaphysical Essay", in U. J. Jensen and Rom Harre (eds.), *The Philosophy of Evolution*. Brighton, U.K.: Harvester Press, 23–44.

Maynard Smith, John, and Eors Szathmáry (1995), *The Major Transitions in Evolution*. Oxford: Oxford University Press.

Michod, Richard (1997), "Cooperation and Conflict in the Evolution of Individuality I. Multilevel Selection of the Organism", *American Naturalist* 149: 607–45.

——— (1999), *Darwinian Dynamics*. Princeton, NJ: Princeton University Press.

Michod, Richard, and Aurora Nedelcu (2003), "On the Reorganization of Fitness during Evolutionary Transitions in Individuality", *Integrative and Comparative Biology* 43: 64–73.

Michod, Richard, and Denis Roze (1999), "Cooperation and Conflict in the Evolution of Individuality III. Transitions in the Unit of Fitness", in Chrystopher Nehaniv (ed.), *Computational Morphogenesis, Hierarchical Complexity and Digital Evolution*, Lectures on Mathematics in the Life Sciences, 26. Providence, RI: American Mathematical Society.

Okasha, Samir (2001), "Why Won't the Group Selection Controversy Go Away?', *British Journal for the Philosophy of Science* 52: 25–50.

——— (2004),"Multilevel Selection and the Partitioning of Covariance: A Comparison of Three Approaches", *Evolution* 58: 486–494.

Roze, Denis, and Richard Michod (2001), "Mutation, Multilevel Selection and the Evolution of Propagule Size during the Origin of Multicellularity", *American Naturalist* 158: 638–654.

Sober, Elliott (1984), *The Nature of Selection*. Chicago: University of Chicago Press.

Szathmáry, Eors, and Lasztity Demeter (1987), "Group Selection of Early Replicators and the Origin of Life", *Journal of Theoretical Biology* 128: 463–486.

Vrba, Elizabeth (1984), "What Is Species Selection?", *Systematic Zoology* 33: 318–328.

https://doi.org/10.1086/508102 Published online by Cambridge University Press

Vrba, Elizabeth, and Stephen Jay Gould (1986), "The Hierarchical Expansion of Sorting and Selection", *Paleobiology* 12: 217–228.
Wilson, David Sloan (1975), "A Theory of Group Selection", *Proceedings of the National Academy of Sciences of the United States of America* 72: 143–146.

https://doi.org/10.1086/508102 Published online by Cambridge University Press