Pluralism, Entwinement, and the Levels of Selection*

Robert A. Wilson†‡

This paper distinguishes and critiques several forms of pluralism about the levels of selection, and introduces a novel way of thinking about the biological properties and processes typically conceptualized in terms of distinct levels. In particular, "levels" should be thought of as being *entwined* or *fused.* Since the pluralism discussed is held by divergent theorists, the argument has implications for many positions in the debate over the units of selection. And since the key points on which the paper turns apply beyond this specific issue, the paper may prove of general interest in thinking about the metaphysics of science

1. Introduction. The "return of the group" has focussed the attention of philosophers and biologists on the relationships among various putative levels at which natural selection operates, particularly on those between group, individual and genic selection. Both unbridled enthusiasts and more circumspect critics of the form of group selection that has received most attention—that of David Sloan Wilson and Elliott Sober (Sober and Wilson 1998; Wilson and Sober 1994)—have endorsed *pluralism* about the

---

*Received June 2002; revised March 2003.

†To contact the author write to Department of Philosophy, 4–115 Humanities, University of Alberta, Edmonton, T6G 2E5, Alberta, Canada; e-mail: rob.wilson@ualberta.ca.

‡An early and brief version of this paper was presented at the Western Canadian Philosophical Association meeting in Edmonton in October 2000, and at the Canadian Philosophical Association meeting in Quebec City in May 2001. I thank my respective commentators, Jim Brown and David Castle, and audiences for their feedback on these occasions. Some of the ideas in the final sections of the paper were informally bruited at a meeting of the Science and Humanities Circle at the University of Alberta in January 2002, and I thank participants there, especially Alex Rueger, for some feedback. Special thanks also to Michael Wade for some valuable written comments on a draft, and for drawing my attention to some of the recent biological literature; and to two referees for *Philosophy of Science,* one of whom drew my attention to the recent work of Kerr and Godfrey-Smith, for their particularly astute comments.

Philosophy of Science, 70 (July 2003) pp. 531–552. 0031-8248/2003/7003-0005$10.00

units of selection. In fact, one or another form of pluralism about the units of selection constitutes the current orthodoxy in the philosophy of biology. The two goals of this paper are to critically examine this pluralistic consensus, and to introduce an alternative to it, one that questions the adequacy of the underlying conception of natural selection as operating at distinct "levels".

I shall distinguish several forms that such pluralism has taken (Section 2), elaborate on the most widely endorsed form of pluralism and the central argument given for it (Section 3), and then critique this argument (Section 4). Pluralism is motivated by the intuition that there may not be a determinate answer to the question of just which level is *the* level at which selection occurs in any particular case. I propose a way to spell out this intuition that breaks from the pluralist consensus (Section 5).

The key positive idea is that various levels of selection are often *entwined* or *fused,* not just in the sense that they co-occur, or operate in the same direction, but in that they are reliably coinstantiated and do not make isolatable, distinct contributions to the ultimate evolutionary currency, fitness. This idea calls into question the view that the levels at which selection operates are sufficiently separable for it to make sense to invoke criteria for preferring one level to others, either in particular cases or in general. I want to suggest not so much the epistemic point that it is sometimes difficult to determine whether individual or group selection is the level at which selection is operating in a particular case as to undermine the claim that there is always a determinate answer to this question. Life, unfortunately, isn't that simple.

This challenge to the idea that the biological world is neatly segmented into layers or levels has broader implications for how we might think about the metaphysics of science. Talk of the various levels at which objects, events, properties, states, and processes exist or can be described is pervasive in discussions within the general philosophy of science and the philosophy of particular sciences (e.g., of physics, biology, psychology, economics). The metaphorical status of such talk has seldom been recognized. While there are reasons to expect entwinement to be pervasive in the biological world, and so for thinking "entwinement talk" to be more revealing than simple "levels talk" in biology, the argument of this paper has implications for these other areas of the philosophy of science. In closing I briefly raise three issues for anyone wishing to explore the notion of entwinement more fully (Section 6).

2. Two Kinds of Pluralism about the Levels of Selection. Consider the label "pluralism". The multilevel theory of selection proposed by Wilson and Sober has been called pluralistic, and it may pay to begin with the contrast between two forms that their own pluralism takes (cf. also Griffiths 2002; Sober and Wilson 1998, 331).

The first and most obvious sense in which Wilson and Sober's view is pluralistic is that it denies that there is any single level at which selection operates; rather, selection can, and often does, operate at multiple levels, levels that often pull in opposite directions, with no one of these levels in general trumping the others. That is, there is a plurality of units of selection—gene, individual, group—with no single one being more fundamental in general than any of the others. As such, this form of pluralism suggests *realism* about the existence of the unit(s) of selection: there are, independent of our particular perspectives, distinct units of selection in the natural world, with their own distinctive properties and subject to particular processes. And it is incompatible with what we might call *fundamentalism,* the idea that one unit of selection is more fundamental (theoretically or ontologically), than the others. I shall call this form of pluralism *unit pluralism,* since at its core is the idea that there is a plurality of units of selection in the biological world itself.

A second and less obvious form of pluralism also exists in Wilson and Sober's multilevel selection theory. It arises in their discussion of the relationship between the theory of group selection and a variety of supposed alternatives to it. According to Sober and Wilson,

> In science as in everyday life, it often helps to view complex problems from different perspectives. Inclusive fitness theory, evolutionary game theory, and selfish gene theory function this way in evolutionary biology. They are not regarded as competing theories that invoke different processes, such that one can be right and the others wrong. They are simply different ways of looking at the same world. When one theory achieves an insight by virtue of its perspective, the same insight can usually be explained in retrospect by the other theories. (1998, 98)

Wilson and Sober go on to argue that the exclusion of the theory of group selection from this "happy pluralistic family" of alternative perspectives "reflects a massive confusion between process and perspective. The theories that were launched as alternatives to group selection are merely different ways of looking at evolution in group structured populations" (1998, 98).

This version of pluralism holds that various prima facie distinct models of natural selection—for example, selfish gene theory and group selection theory—are actually non-competing accounts of one and the same process. I shall refer to this view as *model pluralism,* since it identifies a plurality in the models that evolutionary biologists adopt, rather than in the reality that those models depict. Model pluralism holds that while there may be strategic or pragmatic advantages to using one rather than another model in a particular case, these models do not compete for the truth about the nature of natural selection.

Here I am particularly interested in model pluralism, for this form of pluralism has also been advocated by those with a more sanguine view of the return of the group, such as Sterelny (1996) and Dugatkin and Reeve (1994). Model pluralism has also been elegantly articulated and defended more recently by Kerr and Godfrey-Smith (2002a). According to these authors, there is no fact of the matter as to which of two (or more) putative vehicles or interactors is the level at which selection operates. In particular, descriptions cast in terms of groups and group selection are equivalent to those cast in terms of individuals and what Dugatkin and Reeve call *broad-sense individualism,* the view that "most evolution arises from selfish reproductive competition among individuals within a breeding population" (1994, 107).

Model pluralism has proven popular recently amongst biologists working on the levels of selection (Dugatkin and Reeve 1994, Grafen 1984, Reeve 2000, Dugatkin 2002), the social insects and the origins of eusociality (Bourke and Franks 1995, Crozier and Pamilo 1996), and the evolution of multicellular and social life from simpler forms (Frank 1998, Michod 1999). In general terms, these theorists adopt the view that models positing higher-level processes, such as group selection, do not differ significantly or fundamentally from models positing lower-level processes, such as genic or kin selection. As Bourke and Franks say in summarizing a chapter devoted to this topic, " . . . colony-level, group, individual, and kin selection are all aspects of gene selection. This means that the practice of attributing traits to, say, either colony-level selection or kin selection is illogical" (1995, 67). As this quotation illustrates, model pluralism is sometimes combined with the denial of unit pluralism in that it suggests that there is a sense in which one of these levels—typically, that of the individual or the gene—is more fundamental than the others.

On Sterelny's construal of broad-sense individualism, groups simply form part of the social environment of individuals (cf. Dawkins 1989, 258; Sterelny and Kitcher 1988). Here the form that individual selection takes is *frequency-dependent* selection, where the fitness of a given trait varies as a function of the frequency of either that trait or some other trait in the group or population. Sterelny (1996, 570–573) suggests that this form of pluralism is applicable to a range of cases, including ant colonies/nests and the elaborate chemical and behavioral warning systems that its members have, kin selection in general, honey bee altruism, and trematode altruism (from Dawkins 1982).

Positions similar to model pluralism have been endorsed more generally in debates over the nature of the biological world, particularly when those debates have seemingly either reached a stalement, or when defended positions have led to clashes with widely shared intuitions. For example, faced with the many definitions of biological species, some have concluded that species are “unreal”, or that there are simply different species concepts applicable to different explanatory contexts or projects. And to take an example closer to the model pluralism that is my focus here, Dawkins (1982, ch.1) famously replied to the putative departure that his selfish gene theory made from traditional, individual-centered Darwinian theory by suggesting that the two views were simply different but non-competing perspectives that one could adopt to the unit of selection (see also Dawkins 1989, preface). Shifting between the gene’s-eye view and the individual’s-eye view was like undergoing a gestalt shift in looking at a Necker cube: two views each equally anchored in the reality they reflect (cf. also Grafen 1984 on inclusive fitness). This *Neckerphilia,* as we might call it, has been enthusiastically endorsed most recently by Kitcher (2002) and Kerr and Godfrey-Smith (2002a).

I am suspicious of these ways of deflating prima facie ontological disagreement (see also Wilson 1996, 1999 on species). I shall focus my critique on the model pluralism of Dugatkin and Reeve (1994) and Kerr and Godfrey-Smith (2002a) and the key argument they give for it. But first I lay out the model pluralist view more fully.

3. A Closer Look at Model Pluralism. I have characterized unit pluralism as realist and anti-fundamentalist about the units of selection. By contrast, model pluralism is either conventionalist or fundamentalist about these, views I find objectionable. Dugatkin and Reeve express conventionalism in claiming that “individual and trait-group selection are not alternative evolutionary mechanisms; rather, they are alternative pictures of the same underlying mechanism” (1994, 108). As Sterelny says in glossing Dugatkin and Reeve’s broad-sense individualism, on this view there are only “heuristic differences between the two approaches [broad-sense individualism and group selection theory] . . . the vehicles we should recognize depends on our explanatory and predictive interests” (Sterelny 1996, 572).

Central to the argument of both Dugatkin and Reeve (1994) and Kerr and Godfrey-Smith (2002a) is the claim that there are mathematical, logical, and explanatory equivalences between broad-sense individualism and multilevel selection theory. These are taken to imply that while one might adopt one or the other of these views of natural selection for pragmatic or heuristic reasons, they do not correspond to alternative processes in nature. Rather, the labels “broad-sense individual selection” and “group selection” pick out one and the same process of natural selection, and function in intertranslatable frameworks. Hence, the differences here between broad-sense individual and group selection perspectives are in the eye of the beholder, not in the reality they purport to describe. As Reeve says elsewhere,

> the 'new' group selection models . . . are not mathematically different from the broad-sense individual selection (e.g., inclusive fitness) models at all. Rather, the group selection models are generated from a fitness-accounting scheme that merely produces an alternative *picture* of the same selective processes described by the inclusive fitness models. (Reeve 2000, 65–66)

Much the same view is expressed by Sober and Wilson in the passage quoted in the previous section.

Sterelny's form of model pluralism shifts between this type of conventionalism and fundamentalism. Sterelny's view is fundamentalist in that it claims not only that groups can be regarded simply as parts of an individual's environment, and so group selection is recast as frequency-dependent, individual selection, but also that this implies that the latter is more fundamental than the former. Such fundamentalism derives from Sterelny's skepticism about the status of trait groups in general as units of selection. While he holds that trait groups that form superorganisms are (objectively) vehicles, those that are mere trait groups, i.e., the vast majority of them, are not vehicles (Sterelny 1996, 583) and so there can be no level of selection that operates on them. Thus, a translation scheme that allows one to move from descriptions of group selection to those of individual selection should be taken to imply the eliminability of the former in favor of the latter. With the exception of organism-like groups, trait groups do not form a part of the objective fabric of the biological world subject to natural selection. So alternative descriptions of the level at which selection operates, such as that provided by frequency-dependent, individual selection, are to be preferred. This skepticism about the reality of groups as "vehicles" or "interactors" is widespread, particularly amongst biologists who adopt model pluralism, and introduces a crucial asymmetry between individual (and genic) selection and group selection.

Common to both of these versions of model pluralism is the claim that there are empirically adequate models of evolutionary change that posit different, alternative units of selection. But to derive the specific forms of model pluralism outlined above we need one or the other of two additional assumptions. For Dugatkin and Reeve, this assumption is one about the translatability of claims made in either framework into those of the other; they take this to imply conventionalism about the unit or units of selection. For Sterelny, the missing assumption concerns the unreality of trait groups in general; for him this implies fundamentalism about lower-level units of selection (perhaps excepting the special case of organism-like groups). Although the shared assumption itself warrants further discussion, as does the view of groups that Sterelny articulates, I shall concentrate primarily on the translatability assumption and its putative significance for model pluralism. Consider this line of reasoning, as found in Dugatkin and Reeve (1994) and Kerr and Godfrey-Smith (2002a).

Dugatkin and Reeve claim that "the mathematics of the gene-, individual-, kin- and new group-selection approaches are equivalent" and that "this must be the case and that individual and trait-group selection are not alternative evolutionary mechanisms; rather, they are alternative pictures of the same underlying mechanism" (1994, 108). Their claims of the "logical equivalence" and "logical translatability of one picture into the other" (108) are accompanied by an acknowledgement of the heuristic differences between individual and trait group selection. They take these claims together to entail that claims that pit these two forms of selection against one another are "empirically empty" (109).

Kerr and Godfrey-Smith reinforce this view by showing how the basic parameters in what they call the *individualist* or *contextual* approach and the *multi-level* or *collective* approach to "dealing with group structure in a population" (2002a, 481) can be defined in terms of one another. This allows them to show the parallels between the fitness structures that each of these models generates, and, ultimately, to derive a dynamically sufficient representation of phenotypic change using either model (2002a, 482–484). Since they emphasize the virtues of being able to move between both models, they call their view *gestalt-switching pluralism.* They are Necker-philiacs.

Underlying this sort of argument is the idea that if alternative models of natural selection can be represented either in a common mathematical framework or by the very same equations, then these models differ at most heuristically, not on some deeper level. Dugatkin and Reeve illustrate this by showing how to translate between equations that state when a trait of an individual will evolve by natural selection, and when an allele will be naturally selected. A shared mathematical structure provides the basis for translating between these different models of selection. As they conclude,

> *If broad-sense individual selection, genic selection, and trait-group selection all can be represented by a single condition based only on allele frequencies, then they cannot fundamentally differ from one another* (1994, 109, their emphasis; cf. also Bourke and Franks 1995, 45–49, Box 2.1).

This conditional, however, is false, and the more general argument that it reflects, fallacious. The sense in which various models and forms of selection can be represented by a shared framework, and the significance of this, needs to be explored.

4. Representing, Translation, and Levels. To begin, note that having a shared mathematical framework, or being represented by the very same equations, does not itself entail that two or more processes "cannot fundamentally differ from one another." This is because mathematical models capture just some aspects of the dynamics or kinematics of the processes they model, and they serve *as* models of those aspects only given further assumptions not represented in the models themselves. Indeed, one should expect at least some mathematical models to be applicable to fundamentally *different* processes. Two processes represented in the same mathematical framework or by the same mathematical equations can fundamentally differ from one another because the mathematics (i) is an attempt to selectively represent just some aspects of the phenomenon being modeled, and (ii) requires further, interpretative assumptions in order to serve as a model of that phenomenon.

Consider some examples in areas other than biology. Decision theory is concerned with how agents ought to make choices, and its formal side provides various rules that constrain such choices, including those specified by utility theory, Bayes's account of probabilistic reasoning, and game theory. Such theories can be applied to a range of agents, including human agents, organizations, insect societies, or anything that can be construed as having preferences and means of attaining them. Applying them in each of these cases will require different auxiliary assumptions, but this "pluralism" implies nothing about whether these various domains of application differ fundamentally or not. The backpropagation algorithm familiar from connectionist cognitive modeling in cognitive science can be used to model not only various cognitive processes (verb acquisition, memory, attention, character recognition), but *non-cognitive* systems that can be construed as networks with basic units whose activation is a function of both past and desired levels of activation. Its application within the cognitive realm does not show that the various cognitive processes "cannot fundamentally differ from one another." Likewise, its application to cognitive and non-cognitive domains carries no implication concerning whether these are fundamentally similar or not. Cellular automata have also been used to model cognition as well as the operation of the immune system (Bernardes and Zorzenon Dos Santos 1997), and again the existence of shared mathematical frameworks across domains does little to show just how similar those domains are.

This objection—to the general inference suggested by Dugatkin and Reeve and by model pluralists more generally—does not take us very far, however, for two reasons. First, it begs one question at issue by neglecting (or denying) the *unifying* potential that shared formalisms introduce. Showing how to apply, say, game theory to previously disparate phenomena itself allows those phenomena to be at least partially viewed as in-

TABLE 1. Portion of Kerr and Godfrey-Smith's Table 1 showing fitness structures of the individualist and multi-level perspectives and their relations.

	Individualist	Multi-Level
Definition of fitness structure	α_i = number of copies for an A type in group with *i* A types	π_i = total number of copies for a group with *i* A types
	β_i = number of copies for a B type in group with *i* A types	ϕ_i = number of A copies in a group with *i* A types / total number of copies in a group with *i* A types
Relation to other perspective	$\alpha_i = \pi_i \phi_i / i$	$\pi_i = i\, \alpha_i + (n - i)\, \beta_i$
	$\beta_i = (1 - \phi_i)\, \pi_i / n - i$	$\phi_i = i\, \alpha_i / i\, \alpha_i + (n - i)\, \beta_i$

Reproduced with permission of the authors and Kluwer Academic Publishers.

stances of the same kind of thing; this itself may constitute a major integrative advance. Moreover, the selective nature of mathematical representations is no barrier to their capturing what is central about a phenomenon. One might point to the "frameshifting" between genic, individual, and group selection at the heart of Sober and Wilson's multilevel view of selection as illustrative of both points.

Second, and more importantly, model pluralists about selection have putatively identified an extensive, fine-grained isomorphism between various models of selection. Model pluralists do not simply claim that the models are result equivalent but that they are *causally* equivalent in their structures. For example, Dugatkin and Reeve (1994, 120, Table 1) provide a mapping between individualist and trait-group explanations, and Kerr and Godfrey-Smith claim that their "two parameterizations of the fitness structure in the system contain exactly the same information, parceled up differently" (2002a, 484). Minimally, the onus is on someone who rejects model pluralism to identify differences between the models that cannot reasonably be called "pragmatic" or "heuristic" in nature. I focus below on the parameterizations that Kerr and Godfrey-Smith provide. Their more precise characterization of many of the ideas expressed by Dugatkin, Reeve, and Sterelny reveals the depth of the formal equivalence between "individualist" and "multi-level" perspectives on selection and, in my view, makes the strongest case for model realism yet provided.[1]

Consider Kerr and Godfrey-Smith's Table 1 (2002a, 487), part of which is redrawn above. The fitness structure posited in the individualist per-

---

[1] Their paper appears with commentaries by Dugatkin, Maynard Smith, and Sober and Wilson, as well as their reply to these (Godfrey-Smith and Kerr 2002) and a companion paper that continues their argument by focusing on models that use Price's equation and average fitness (Kerr and Godfrey-Smith 2002b).

spective is parameterized in terms of the number of copies of the relevant genotypes or phenotypes, α_i and β_i, where the index signifies a relativization to a group with *i* A types. This reflects a standard individualist view, adapted so as to explicitly recognize a role for population structure as the *context* in which types occur. The fitness structure in the multi-level perspective, by contrast, is parameterized in such a way as to reflect the group as a *collective* that itself can serve as the bearer of fitness values, and so captures the idea that natural selection can operate on groups themselves. Thus, π_i and ϕ_i aim to express, respectively, group productivity and the contribution of A types to that productivity. So parameterized, the mathematical equivalences between the two perspectives falls out with some simple algebra.

Suppose that we grant the parameterizations for now. What follows? To derive a conclusion about equivalence in the two *perspectives,* a conclusion about the translatability between them, we need at least the assumption that these formalizations express the core of those perspectives. One reason to question this, even granting the formalizations, is that the two perspectives differ in their quantificational range. In particular, the individualist perspective quantifies over just *individuals* while the multi-level perspective quantifies over *groups,* not only as the bearers of fitness values or interactors but also as the manifestors of adaptations and the beneficiaries of natural selection, to use Elisabeth Lloyd's (2001) useful terminology. A group may play one or more of these roles in any given case in the multi-level perspective, but *none* of these roles in the individual perspective. Groups, in short, are *agents* in the process of natural selection in the multi-level perspective, but are never such in the individualist perspective. This is an ontological difference between the two perspectives, and insofar as it is not reflected in the formalizations, they leave something important out. Moreover, it is difficult to see how adding in parameters to the multi-level model to capture such agentive roles would allow one to preserve the isomorphism between the two formal models, since that would in turn require corresponding parameters in the individualist model. But since individualists do not think that groups play such roles—certainly, not in *all* cases (see Maynard Smith 2002)—this would seemingly require a more radical departure from their perspective than simply relativizing types to population structure.

There is a dilemma here for the model pluralist that we might express in several ways. Consider one of these. If the models are exhaustively constituted by the formal definitions provided by Table 1 (and its extensions), then these models leave out something crucial about each of the perspectives they represent. So formal equivalences in the models do not support the claim that the two perspectives are simply different ways of expressing a shared view of the biological world (and so do not compete for the truth about that world). If, on the other hand, we need to extend the formalizations of each perspective so as to reflect the deeper ontological differences between them, then the onus is on the model pluralist to show how to (a) perform the formalization, and (b) preserve the isomorphism between the two formal models. I have no definitive argument to show that this is impossible, though I have said something briefly about why I suspect that it cannot be done.

Let us return to the parameterization itself, for one of its interesting features is that it provides a way to translate formally between the individualist and multi-level views of the bearer of fitness values, interactors. If this is successful, then it captures one of the "ontological" differences between the perspectives—whether groups play this role—itself a significant achievement. But the characterization of both parameters, π and ϕ, should be questioned by proponents of the multi-level perspective, precisely because each ties a group's productivity too closely to the expected numbers of copies of the genotypes or phenotypes in that group. In some circumstances a group's productivity may be measured just by the total number of copies of individuals it leaves, as the parameterization suggests, but such a measure will be misleading in others where the group structure itself plays a significant role in directing evolutionary change.

For example, consider a group, T, with A and B types that reproduces three groups, U, V, and W, each of its original size, as depicted in Figure 1a. At generations 1 and 2, this lineage has precisely the same group productivity (fitness) as defined by π as the lineage in Figure 1b. Yet since group size and structure may be ecologically significant, there may be a critical fitness difference between the two lineages, as revealed in generation 3. For example, the T lineage contains Z, which will allow the replication pattern in generation 2 to be reproduced again. And in the Q lineage, R may exceed a threshold, leading to a lower fitness level. Moreover, differences between the groups in Figure 1a after the first generation may themselves be significant for evolutionary change. For example, a "pure group" like W breaks the dynamic of decreasing within-group fitness of As relative to Bs. Note that this dynamic is broken in the transition from T to W, and from V to Z, through the "defection" of just one B to A, something that itself may be a function of group structure. π is insensitive to these sorts of *causal* differences precisely because it conceptualizes group productivity solely in terms of the composition of groups rather than the reproduction and structure of groups themselves.

One might grant that this is true but insignificant because it is the other parameter, ϕ that is relevant here, for it is here that the underlying causal story is to be told within the multi-level perspective. Again, however, this should be questioned by a proponent of the multi-level perspective, for this parameterization is also cast exclusively in terms of *individual* off-

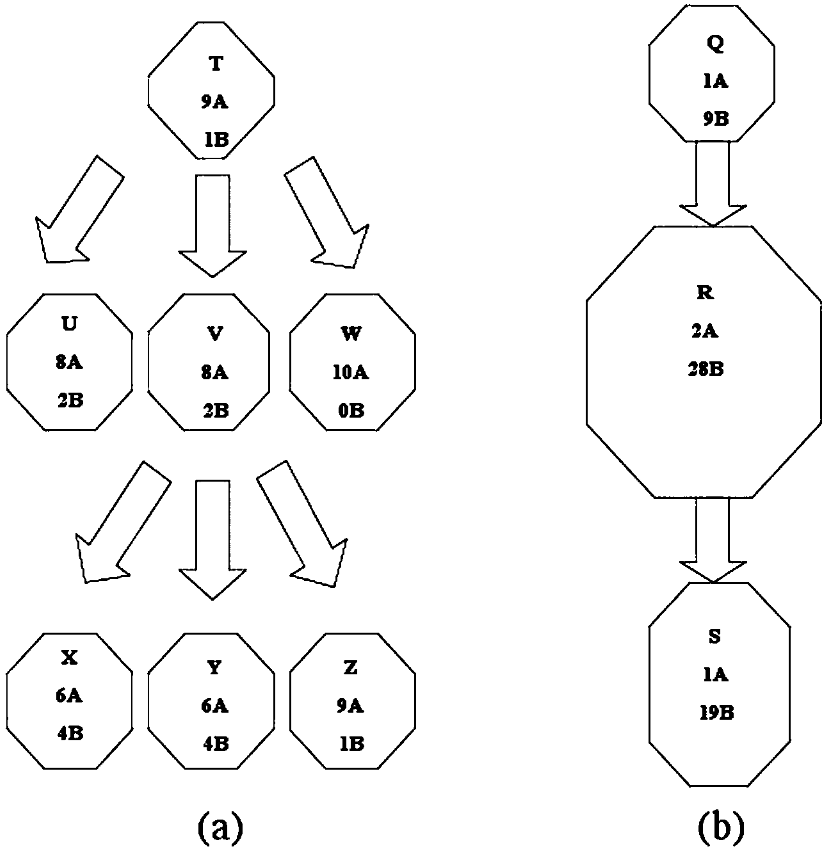

Figure 1. (a) The T lineage, containing A and B types. The portion of generation 3 depicted shows how the group structure iterates through group Z. (b) The Q lineage, whose group R in generation 2 exceeds some threshold, leading to a fitness drop. In general, B increases in relative fitness across generations, with the exceptions (from T to W, and from V to Z) mediated by the group structure in the T lineage.

spring or copies. But just as group productivity should sometimes be understood within the multi-level perspective in terms of the production *of groups* (and not just individuals in those groups), we also need to look beyond an individual's within-group frequency in order to understand the causal interplay between individuals and their population-structured environments. As Figure 1 illustrates, having a certain proportion of A types within a group may be necessary for that group to flourish or even survive, and this role may vary across different ecological circumstances. In such a case, the parameterization will offer at best a partial characterization or approximation of the contribution of A types to group productivity.

The basic problem with both parameters is precisely what allows them to serve as the bridge between the individualist and multi-level perspectives: they are themselves too individualistic. Group productivity is parameterized solely in terms of the individual composition of groups and then used to define the particular contributions of individuals to overall group productivity. While this is sometimes an appropriate way to understand groups themselves as bearers of fitness values—since group fitness sometimes *just is* the sum of the fitness of the individuals it produces—it does not allow one to represent circumstances in which groups play a more agent-like role in natural selection.

One final point about model pluralism. I noted at the outset that model pluralism has proven attractive across the spectrum of views over the levels of selection. But this attraction has a very different basis for proponents of group selection and individualists about selection. David Sloan Wilson, for example, has sought to restore the theory of group selection to respectability over an extended period in which individualistic views have predominated. Thus, he has both distanced his views from "naive group selection theory" of the past and sought to integrate his "new group selection theory" with predominant individualistic views. In this context, model pluralism serves a legitimating function. Individualists, by contrast, are attracted by model pluralism because it allows them to treat new group selection theory as old wine in new bottles. Cases that seem to call out for viewing groups themselves as occupying centre-stage in the process of natural selection are redescribed from the individualist perspective with groups present but very much in the wings. Phenomena that feature groups themselves as units of selection exist but are rare (cf. Sterelny 1996, Maynard Smith 2002), very much the view that Williams (1966) originally defended. In this context, model pluralism functions to maintain the status quo.

Disagreements that lie beneath the surface of the pluralist consensus reflect this contextual difference. One such disagreement concerns unit pluralism and its realist and anti-fundamentalist standpoint (see also Sober and Wilson 2002). We have seen that Sober and Wilson endorse both model pluralism and unit pluralism, while individualists remain tempted by fundamentalism and are, at best, neutral or guarded about unit pluralism. One way to read my criticism of the translatability argument for model pluralism is as suggesting that unless model pluralism allows us to resolve the *ontological* issues that unit pluralism raises, it stops short of dissolving the disagreement between individualists and multi-selectionists.

5. Entwined Levels of Selection. Model pluralism is an enticing position for making sense of debate over the units of selection, a temptation harder to resist in light of the appeals to the "mathematical equivalences" between various models of selection. But I have argued that these equivalences do not have the significance that model pluralists have ascribed to them. And I view model pluralism itself as a mistaken way to spell out the intuition that there may not be a determinate answer to the question of just which level is “the” level at which selection occurs in any particular case. Unit pluralism shares this intuition but adopts a realist view of whether selection is operating at the genic, individual, or group level (or all three) in any particular case. The unit pluralist’s supposition is that the levels at which selection operates are always sufficiently separable for it to make sense to invoke criteria for preferring one or more levels to others, at least in particular cases.

It is this supposition that I challenge in this section. Various levels of selection are often *entwined* in that while they are reliably coinstantiated, they do not make isolatable, distinct contributions to the ultimate evolutionary currency, fitness. Especially when two or more levels of selection work in the same direction, and act in ways to mutually reinforce one another, as I think is often the case, isolating the contribution of each level involves an artificial partitioning of the causes of the resulting evolutionary change (or stasis) that does not track biological reality. In what follows I hope to show how to articulate this idea to arrive at a position that departs from both unit and model pluralism while accommodating the central intuition that motivates these views.

I present the argument in two stages: with respect to properties *at a level,* and then with respect to processes or mechanisms *at different levels.* The first stage aims to show that properties at the group level can be entwined, as can those at the individual level. In the second stage I argue that the very processes that generate these entwined properties across the two levels are themselves entwined. It is this entwinement “across levels” that challenges the very idea that the biological world separates neatly into layers or levels. I develop both stages of the argument through a detailed, idealized example.

Consider a metapopulation consisting of groups of nesting organisms that vary with respect to their invadability by outsiders. Suppose that the least invadable groups have a certain proportion of distinctly larger individuals, and this is because those individuals are particularly adept at blocking holes in the nest with their larger-than-average bodies. These larger individuals come to specialize in nest defense, and there is some division of labor within the nest of which such a specialization forms a part. Since larger individuals require more food resources, too many of these lead to resource-based problems for the group, and so there is an equilibrium number, *m,* of large individuals for any nest of size *n* to possess. Here it would be natural to say that certain groups are less invadable *because* they have a certain proportion, *m/n,* of larger individuals. Suppose also that individuals are able to grow distinctly larger only because they are members of relatively uninvadable groups, since the corresponding division of labor creates time for more individuals to forage. Suppose also, as a result, that uninvadable groups have a higher level of fitness than invadable groups.

In a population with very few distinctly large individuals there will be individual selection for individuals of greater body size: their proportion will increase within the group over time until it reaches the equilibrium ratio of *m*/*n*. This is a form of frequency-dependent selection that selects for an intrinsic property of individuals, body size. But since all and only those individuals with greater body size are also those who function as hole-blockers, there is also selection for hole-blockers in this population. In the environment in which selection actually occurs the realization of one of them is also a realization of the other, and this is no accident. Natural selection is sensitive to the ways in which properties are reliably instantiated, and so it is not fine-grained enough to distinguish between entwined properties.

Consider now group selection. Because the groups vary with respect to their invadability, this variation correlates with fitness variation, and that variation is heritable (i.e., survives intergenerationally), lower levels of invadability are adaptations and there is a (group) selective pressure for their evolution. Either groups with low levels of invadability will differentially grow in members, or such groups will be ancestors to more groups than those with higher levels of invadability. But it is also surely true that there is group selection for groups with, or that approximate having, *m*/*n* large individuals. Being uninvadable and having *m*/*n* large individuals are not identical—there are possible, even nomologically possible, situations in which they come apart—yet in the example as I have described it they are entwined. And as before, natural selection does not—indeed, *cannot*—distinguish between entwined properties.[2]

Furthermore, not only are properties at both the group-level and the individual-level entwined, but the two corresponding processes or mechanisms, group and individual selection, are also entwined. Individuals with large body size and that play the hole blocker function are selected within groups, and there is the selection between groups of lower levels of invadability and having *m*/*n* large individuals. But the selective processes at different levels here are not independent; they constitute mutually rein-

---

[2] Sober's (1984) distinction between the selection *for* and the selection *of* properties is irrelevant to both cases of entwinement I have sketched here. Some will find this easier to recognize if the descriptions are cast in terms of there being *laws* that relate the entwined properties. Since I think that an appeal to biological laws, especially in this context, is misplaced, I talk instead of the properties being reliably coinstantiated.

forcing forces whose effects, the properties that are selected, are closely and deeply related, with neither individual- nor group-level properties being causally prior to the other. These reliably coinstantiated selective forces do not make isolatable, distinct contributions to fitness. Moreover, given the case as described, these selective processes can be teased apart only in thought or through experimental interventions that alter the very relationships between them.

Perhaps what I am calling entwinement would be more aptly named *fusion,* since its occurrence implies the oneness of what might, in other cases, be separate and separable forces, in much the way that two winds can fuse to create a gale whose effects are not attributable separably to each of the initially (and in other circumstances, continuingly) distinct winds. Entwinement or fusion of properties or processes implies that it makes no sense to apportion determinate, partial causal responsibility for the resulting evolutionary change. This is true whether we take the properties of having a certain body size or type and being a hole-blocker at the individual level, of being uninvadable and having *m*/*n* large individuals at the group level, or the levels themselves.

When either properties or processes *merely* co-occur, they can be teased apart through experimental means, using various control groups and experimental conditions in order to isolate the distinctive contribution that a property or process makes to the direction and size of the force of natural selection. This is precisely what one might think lies at the heart of experimental approaches to determine the level or levels at which selection operates. But such experimental interventions reach a limit when properties or processes are entwined, since in that case the two do not pre-exist as independent, separable parts of the world. When there is entwinement, further empirical investigation reveals further complexities and entanglements, and the discovery of the workings of isolated properties or processes is an artifact of the experimental investigation, an artifice of our imposing on the world an ontology with distinct levels in an idealized biological hierarchy.

An examination of Michael Wade's (1977, 1978, 1980, 1982) classic experimental approach to group selection on flour beetles, two species of *Tribolium,* and the work that it has inspired, will illustrate this point. In a series of elegant and often-discussed experiments, Wade and his colleagues selected for groups of beetles with high rates of population increase, pitting these against both randomized control groups and those in which there was only individual selection. Rather than using a method of panmixis across generation times, as had previous models of group selection, Wade used propagules as the basis for new generations in his artificial (group) selection model. The basic finding was that there are conditions under which the effects of group selection are significant, and significantly stronger than those of individual selection. Wade himself selected for the group-level property of rate of population increase, but related studies have selected for other group-level traits in *Tribolium,* such as emigration rate (e.g., Craig 1982), and multi-level traits in other species, such as leaf area in the cress *Arabidopsis thaliana* (Goodnight 1985).

Although each of these studies shows that a certain property can be most effectively selected through group rather than individual selection, even collectively they do not show that there are *no* other properties being selected through either group or individual selection. Goodnight and Stevens have noted that "[a]ll of the traits examined in group selection experiments are influenced by interactions among individuals" (1997, S67): population size and emigration rates (in *Tribolium*) by cannibalism rates and population density, respectively; and leaf area (in *Arabidopsis*) by planting density. Each of these pairs of properties is entwined in the experimental conditions, and thus both are selected through group selection (high population size with low rates of cannibalism; high emigration rates with high population density). But there is also a range of individual-level, dispositional properties, such as fecundity, aggression, and cannibalism, that are not independent of these group-only and multi-level traits, and that are selected along with them. For example, when the rate of cannibalism is high, there are more individuals manifesting that trait, and so we could see individual selection as operating on *it.* This sort of entwinement underpins Sterelny's suggestion that we "reanalyse the evolution of (say) large flour beetle groups as the evolution of a gene for (say) synchronizing breeding, a gene advantageous only in a particular population-structured context" (1996, 571). What I am suggesting is that we give up the fundamentalism that I think drives Sterelny's own talk of "reanalyzing", but recognize the reality of the metaphysical entwinement that gives rise to it.

Wade's more recent work on "indirect genetic effects" (IGEs) is interestingly positioned vis-à-vis traditional talk of levels of selection and the notion of entwinement that I have introduced. IGEs occur when there are phenotypic or fitness changes that result from changes to the genes of an individual's social partners, rather than simply from that individual's own genes. While phenotypic expression and fitness always depend on the values of environmental variables, IGEs constitute a malleable environment whose effects on selection can be extreme. For example, the effect of plant height on a given plant's fitness will depend, in part, on the heights of its surrounding neighbors, and changes in an individual's genetic environment can very quickly change the fitness values of its phenotypes. Agrawal, Brodie III, and Wade (2001) extended the basic models of IGEs to apply to cases in which population structure directs selection, finding that

IGEs could be significantly higher once assumptions of panmixis and linear interactions were relaxed. They remind us (2001, 318) that the non-additive effects of individual and group selection had long been noted in the literature on experimental approaches to group selection, and view their own models as providing an account of the interactions between these levels of selection. This is because the models themselves add in terms that explicitly refer to the effects of interactions between the levels of selection. However, one might also see this work as posing a challenge to the framework that multi-selection theory operates within, and as providing conceptual and mathematical tools more readily viewed as applying not to a layered biological reality but to one whose "levels" are inherently entwined with one another. Wolf, Wade, and Brodie III (2003), which proposes a general way to model a variety of gene-context interactions, might be viewed as a further step in this direction.

The entwinement of properties and processes in the biological world is clearly a matter of degree, a point that we can use to situate the view I am advocating with respect to our two forms of pluralism. When entwinement is extremely strong or tight, as it was in my idealized example, then model pluralism takes on some appeal, since the modeling "tools of the trade", like natural selection itself, will be insensitive to underlying ontological differences. And when entwinement is extremely weak or loose, unit pluralism will allow us to approximate the contributions of individual-level and group-level properties and agents to the process of evolution by natural selection. Since I think that we have significantly under-estimated the extent to which properties in the biological world are entwined, the limitations of unit pluralism are likely to be more severe than its proponents have allowed.

The idea that biological reality is entwined provides a different starting point, however, from both model and unit pluralism in thinking about the nature of selection. For any given organism the evolutionary interest of various units—genes, individuals, and groups—are inextricably woven together. None of these putative units floats free of the others in evolutionary space. To take the standard Darwinian unit, individuals typically mediate their reproductive success through their membership in a group or groups, and do so by means of the replication of their genetic material; genes exist and thus reproduce only in individuals, and groups do so only through the reproductive efforts of the individuals that comprise them (cf. Heisler and Damuth 1987, Damuth and Heisler 1988). There is thus a clustering of replicative, reproductive, and social interests—not always in harmony but inextricably knitted together—that lies somewhere between the merely fortuitous and the lawful. By taking this entwinement to be ontologically basic, we can see the hierarchically-layered view shared by pluralists, and so the issue of how these levels are related, as an artifact of just one way of thinking about the metaphysics of science.

6. Levels, Entwinement, and the Metaphysics of Science. In the previous section I introduced the idea of entwinement as a tool for complicating the conception of natural selection as operating at distinct levels that underlies the forms of pluralism discussed in Sections 2–4. Given the ubiquity of talk of levels amongst those engaged in various debates over the metaphysics of science, if the notion of entwinement proves useful in characterizing the levels of selection, it is likely to have broader application. I want to conclude the substantive part of the paper by raising some general issues about levels and entwinement with this in mind.

The first of these concerns the *metaphorical* nature of talk of levels in many scientific contexts. Such talk might be thought to be literal, to correspond to a distinctly layered reality, in physics, where the concern is to mark scale differences between objects of different size (e.g., Dresden 1998). Whatever one thinks of this, it seems to me too optimistic as a view of levels-speak in other sciences. For example, the common idea that we should understand talk of the "levels of selection" literally in terms of entities nested in a compositional hierarchy (e.g., genes, individuals, families, groups) is confused. First, there is no univocal relational expression that applies to the full range of entities that have been postulated as putative objects of selection. "Compose", "are a part of", "physically constitute", and "are realized in"—to take four of the more obvious predicates relating smaller to larger biological entities that are candidate units of selection—do not truly, literally, and univocally describe the relationship between adjacent pairs in the following series: alleles, genotypes, cells, organs, bodily systems, organisms, kin groups, families, nests, colonies, species. Second, it is much less clear how properties, processes, events, and states of such entities can stand in these sorts of mereological relationships to one another. For example, supposing that groups are physically composed of individuals, in what sense is the *process* of group selection physically composed of individual selection? At best, the relationship between these processes would seem to be derivative, in some way, from that which holds between the physical objects to which they are "attached". This problem is more severe in the philosophy of mind and the cognitive sciences, where the levels are often simply designated "mental" and "physical" or their variants "psychological" and "neurological".

The second of these issues concerns the relationship between a layered and an entwined metaphysical view of scientific ontology. There seem two distinct ways to think of the relationship. On the one hand, one could view entwinement as an *alternative* to the layered view of the world. If one thought that talk of levels was misleading, or was itself responsible for creating some of the puzzles that keep philosophers busy (e.g., the mind-body problem), then this is likely to be an attractive view of the entwinement metaphor. If the entwinement metaphor is developed with this in mind, then talk of *fusion* may prove more apt and useful than that of entwinement. On the other hand, one could view entwinement as a *supplement* to the view of the world as layered, one that either is applicable to special cases and circumstances or that allows one to fine-tune or correct the predominant layered view of scientific ontology. Both of these views may be useful in thinking further about the nature of selection.

The third issue is how we can make the notion of entwinement both more precise and more general, for at root it too gains much of whatever appeal it has to its metaphorical nature. Despite the pioneering work of Mary Hesse (1966) and Richard Boyd (1993), both building on Max Black's (1962) views of metaphor, that established a place for metaphor in science, this remains an uncomfortable place for many offering an analytic understanding of scientific processes and debates over them. I have suggested that phenotypic properties are entwined at a level just when they are reliably coinstantiated but do not make distinct contributions to the fitness of their bearer, and that processes across levels are entwined just when *they* are reliably coinstantiated but do not make distinct or independent contributions to the fitness of their respective bearers. There are natural ways to generalize such notions of entwinement to properties and processes more generally, but how plausible these are will turn, in part, on how we conceptualize the relationship between levels and entwinement.

7. Conclusion. Model pluralists are right that there is, in at least some cases, no fact of the matter as to whether selection operates at the (broad-sense) individual level or at the group level. This is not, however, because these two models are inter-translatable, or because one is reducible to the other, but rather because the reality that these models describe includes entwined or fused properties and processes. Thus the resulting indeterminacy as to just which property is selected, or just which process governs its selection, is not a legacy of our epistemic limitations. Instead, it is a reflection of a biological reality that does not always fall neatly into distinct levels, as proponents of unit pluralism imply. If this introduction of entwinement stimulates discussion in this and other areas in which philosophers of science and others have complacently appealed to the metaphor of levels, so much the better.

REFERENCES

Agrawal, Aneil F., Edmund D. Brodie III, and Michael J. Wade (2001), "On Indirect Genetic Effects in Structured Populations", *American Naturalist* 158: 308–323.

Bernardes, A.T., and R.M. Zorzenon Dos Santos (1997), "Immune Network at the Edge of Chaos", *Journal of Theoretical Biology* 186: 173–187.

Black, Max (1962), *Models and Metaphors.* Ithaca, NY: Cornell University Press.

Bourke, Andrew F.G., and Nigel R. Franks (1995), *Social Evolution in Ants.* Princeton, NJ: Princeton University Press.

Boyd, Richard (1993), "Metaphor and Theory Change: What is 'Metaphor' a Metaphor For?", in Andrew Ortony (ed.), *Metaphor and Thought.* 2d ed. Cambridge: Cambridge University Press, 481–532.

Craig, D. M. (1982), "Group Selection Versus Individual Selection: An Experimental Analysis", *Evolution* 36: 271–282.

Crozier, Ross H., and Pekka Pamilo (1996), *Evolution of Social Insect Colonies.* New York: Oxford University Press.

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

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

——— (1989), *The Selfish Gene.* 2d ed. Oxford: Oxford University Press.

Dresden, Max (1998), "The Klopsteg Memorial Lecture: Fundamentality and Numerical Scales—Diversity and the Structure of Physics", *American Journal of Physics* 66: 468–482.

Dugatkin, Lee Alan (2002), "Will Peace Follow?", *Biology and Philosophy* 17: 519–522.

Dugatkin, Lee Alan, and Hudson K. Reeve (1994), "Behavioral Ecology and Levels of Selection: Dissolving the Group Selection Controversy", in Peter J.B. Slater et al. (eds.), *Advances in the Study of Behavior,* vol. 23. New York: Academic Press, 101–133.

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

Godfrey-Smith, Peter, and Benjamin Kerr (2002), "Group Fitness and Multi-Level Selection: Replies to Commentaries", *Biology and Philosophy* 17: 539–549.

Goodnight, Charles (1985), "The Influence of Environmental Variation on Group and Individual Selection in a Cress", *Evolution* 39: 545–558.

Goodnight, Charles, and Lori Stevens (1997), "Experimental Studies of Group Selection: What Do They Tell Us About Group Selection in Nature?", *American Naturalist* 150: S59-S79.

Grafen, Alan (1984), "Natural Selection, Kin Selection, and Group Selection", in J. Krebs and N. Davies (eds.), *Behavioural Ecology: An Evolutionary Approach.* London: Blackwell, 63–84.

Griffiths, Paul (2002), Review of Sober and Wilson 1998, *Mind* 111: 178–182.

Heisler, I. Lorraine, and John Damuth (1987), "A Method for Analyzing Selection in Hierarchically Structured Populations", *American Naturalist* 130: 582–602.

Hesse, Mary (1966), *Models and Analogies in Science.* South Bend, IN: University of Notre Dame Press.

Kerr, Benjamin, and Peter Godfrey-Smith (2002a), "Individualist and Multi-Level Perspectives on Selection in Structured Populations", *Biology and Philosophy* 17: 477–517.

——— (2002b), "On Price's Equation and Average Fitness", *Biology and Philosophy* 17: 551–565.

Kitcher, Philip (2002), "Evolution Theory and the Social Uses of Biology", paper read at eighteenth biennial meeting of the Philosophy of Science Association, November 2002, Milwaukee, WI.

Lloyd, Elisabeth A. (2001), "Units and Levels of Selection: An Anatomy of the Units of Selection Debates", in Rama S. Singh, Costas B. Krimbas, Diane B. Paul, and John Beatty (eds.), *Thinking About Evolution: Historical, Philosophical, and Political Perspectives.* New York: Cambridge University Press, 267–291.

Maynard Smith, John (1987), "How to Model Evolution", in John Dupré (ed.), *The Latest on the Best.* Cambridge, MA: MIT Press, 119–131.

——— (2002), "Commentary on Kerr and Godfrey-Smith", *Biology and Philosophy* 17: 523–527.

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

Reeve, Hudson K. (2000), "Multi-Level Selection and Human Cooperation", *Evolution and Human Behavior* 21: 65–72.

Sober, Elliott (1984), *The Nature of Selection.* Cambridge, MA: MIT Press.

Sober, Elliott and David Sloan Wilson (1998), *Unto Others: The Evolution and Psychology of Unselfish Behavior.* Cambridge, MA: Harvard University Press.

——— (2002), "Perspectives and Parameterizations: Commentary on Benjamin Kerr and

Peter Godfrey-Smith's 'Individualist and Multi-Level Perspectives on Selection in Structured Populations'", *Biology and Philosophy* 17: 529–537.

Sterelny, Kim (1996), "The Return of the Group", *Philosophy of Science* 63: 562–84.

Sterelny, Kim and Philip Kitcher (1988), "The Return of the Gene", *Journal of Philosophy* 85: 339–361.

Wade, Michael J. (1977), "An Experimental Study of Group Selection", *Evolution* 31: 134–153.

——— (1978), "A Critical Review of the Models of Group Selection", *Quarterly Review of Biology* 53: 101–114.

——— (1980), "An Experimental Study of Kin Selection", *Evolution* 34: 844–855.

——— (1982), "Group Selection: Migration and the Differentiation of Small Populations", *Evolution* 36: 944–961.

Williams, George C. (1966), *Adaptation and Natural Selection.* Princeton, NJ: Princeton University Press.

Wilson, David Sloan, and Elliott Sober (1994), "Reintroducing Group Selection to the Human Behavioral Sciences", *Behavioral and Brain Sciences* 17: 585–654.

Wilson, Robert A. (1996), "Promiscuous Realism", *British Journal for the Philosophy of Science* 47: 303–316.

——— (1999), "Realism, Essence, and Kind: Resuscitating Species Essentialism?", in Robert A. Wilson (ed.), *Species: New Interdisciplinary Essays,* MIT Press, Cambridge, MA, 187–207.

Wolf, Jason B., Michael J. Wade, and Edmund D. Brodie III (2003), "The Genotype-Environment Interaction and Evolution when the Environment Contains Genes", in Thomas J. DeWitt and Samuel.M. Scheiner (eds.), *Phenotypic Plasticity: Functional and Conceptual Approaches.* New York: Oxford University Press.