

MULTILEVEL SELECTION AND THE PARTITIONING OF COVARIANCE: A COMPARISON OF THREE APPROACHES

SAMIR OKASHA

Department of Philosophy, University of Bristol, 9 Woodland Road, Bristol, BS8 1TB, United Kingdom.

E-mail: Samir.Okasha@bristol.ac.uk

Abstract.—Where the evolution of a trait is affected by selection at more than one hierarchical level, it is often useful to compare the magnitude of selection at each level by asking how much of the total evolutionary change is attributable to each level of selection. Three statistical partitioning techniques, each designed to answer this question, are compared, in relation to a simple multilevel selection model in which a trait's evolution is affected by both individual and group selection. None of the three techniques is wholly satisfactory: one implies that group selection can operate even if individual fitness is determined by individual phenotype alone, whereas the other two imply that group selection can operate even if there is no variance in group fitness. This has significant implications both for our understanding of what the term "multilevel selection" means and for the traditional concept of group selection.

Key words.—Altruism, contextual analysis, covariance, group selection, multilevel selection, Price's equation.

Received April 14, 2003. Accepted October 14, 2003.

Multilevel selection theory is the idea that natural selection may occur simultaneously at different levels of the biological hierarchy. This idea has generated considerable interest in the recent literature. Multilevel selection theory has proven useful for analyzing a variety of phenomena, including the evolution of social behaviour (Wilson 1997; Frank 1999), intragenomic conflict (Pomiankowski 1999), the origin of multicellularity (Michod 1997, 1999), and the origin of chromosomes (Maynard Smith and Szathmari 1995).

Where the evolution of a trait, or a gene, is affected by selection at more than one level, it is often useful to ask which level of selection is likely to dominate. This is because the direction of selection may be different at different hierarchical levels; in the classic example, altruistic behavior is usually selected against at the individual level, but selected for at the group level. Similarly, in Michod's models of the evolution of multicellularity, defecting cells, which abandon somatic duties in favor of increased replication, are favored by selection at the cell level but disfavored by selection at the organism level.

From a theoretical perspective, the most general way of comparing the relative strengths of two (or more) levels of selection is to partition the total evolutionary change into a number of components, each representing the effect of selection at a different level. Surprisingly, however, there is no consensus on how exactly this partitioning should be done. In this paper I examine three statistical partitioning techniques, each of which is found in the recent literature on multilevel selection and each of which purports to divide the total change into components corresponding to distinct levels of selection. The three techniques are not equivalent; their interrelationships are briefly explored. I argue that there is no obvious way of choosing between the three, as each has points in its favor. This, in turn, implies that the concept of multilevel selection is less clear-cut than is often assumed.

One preliminary point: for simplicity I assume just two levels of selection, the individual level and the group level. But it should be easy to see how to extend the model below and the statistical partitioning techniques to cover more than two levels, or levels other than individuals and groups.

The Price Approach

The first and best-known partitioning technique stems from the celebrated Price equation. Consider first a single unstructured population, containing no groups. Suppose that a measurable phenotypic character of organisms, denoted by z , is undergoing directional selection, so the mean character in the population, \bar{z} , will change over time. Price (1972) showed that the observed selection differential on the character z , defined as the within-generation change in \bar{z} due to natural selection, is given by:

$$\Delta\bar{z} = \frac{\text{Cov}(w_i, z_i)}{\bar{w}}, \quad (1)$$

where z_i is the phenotypic value of the i th organism, w_i is the absolute fitness of the i th organism, \bar{w} is mean fitness, and Cov denotes the covariance function.

Equation (1) requires that the phenotypic character z be assigned a metric, so it applies in the first instance to continuous characters such as height, rather than discrete characters such as eye color. But this is not a serious limitation. Suppose eye color is the trait undergoing selection, and we are interested in the proportion of brown-eyed organisms in the population. We simply define $z_i = 1$ if the i th organism is brown eyed, and $z_i = 0$ otherwise. Obviously, \bar{z} then equals the proportion of brown-eyed organisms in the population. Similarly, z could be defined as the frequency of a particular allele at a given locus in an organism ($= 1, \frac{1}{2}$, or 0 for diploid species); \bar{z} would then equal the overall frequency of the allele in the population. Equation (1) can therefore handle discrete as well as continuous characters, and genotypic as well as phenotypic ones.

Now we introduce multilevel selection. Suppose the population is subdivided into a number of groups, assumed for simplicity to be of equal size. Nothing in particular about the nature of the groups is assumed—they could be kin groups, colonies, demes, or anything else. The nature of the groups, and the question of their biological justification, is discussed below, in *Group Selection: A Concept in Tension*. We define z_{ij} as the phenotypic value of the i th organism in

the j th group. We define Z_j as the average phenotypic value of the organisms in the j th group, and W_j as the average fitness of the organisms in the j th group. So $Z_j = 1/n \sum_i z_{ij}$ and $W_j = 1/n \sum_i w_{ij}$, where n is group size. We can think of Z_j as the group phenotype or group character of the j th group and W_j as its group fitness.

Since the population is subdivided into groups, it is possible to partition the total selection differential on the character z , given in equation (1), into within-group and between-group components, as Price (1972) showed. This gives

$$\bar{w}\Delta\bar{z} = E_j[\text{Cov}_W(z_{ij}, w_{ij})] + \text{Cov}(Z_j, W_j). \quad (2)$$

The first term on the right side of equation (2), $E_j[\text{Cov}_W(z_{ij}, w_{ij})]$, is the average of the within-group covariances between z and w . The second term, $\text{Cov}(Z_j, W_j)$, is the covariance between group phenotype Z_j and group fitness W_j . Equation (2) involves no new biological assumptions; it follows directly from equation (1), given that the population is subdivided into groups. See Wade (1985) or Price (1972) for a full derivation of equation (2).

Numerous authors have regarded equation (2) as a useful way of modeling the combined effects of group selection and individual selection on a character. The equation appears to partition the total selection differential on a character into two components, corresponding respectively to individual selection and group selection. Price (1972), Hamilton (1975), Wade (1985), Sober and Wilson (1998), and others have all endorsed this suggestion. In their view, $E_j[\text{Cov}_W(z_{ij}, w_{ij})]$ represents the effect of individual selection on the change in \bar{z} , while $\text{Cov}(Z_j, W_j)$ represents the effect of group selection. I call this the ‘‘Price approach’’ to multilevel selection.

The grounds for the Price approach are easy to see. Suppose that all the groups have the same fitness, or that group fitness and group character are uncorrelated. Intuitively there can be no group selection in such a situation. And the term $\text{Cov}(Z_j, W_j)$ will equal zero in that situation. Conversely, suppose that within each group, all organisms have the same fitness, or that organismic fitness and character are uncorrelated within each group. Intuitively there can be no individual selection in such a situation. And the term $E_j[\text{Cov}_W(z_{ij}, w_{ij})]$ will be zero in that situation—for each of the within-group covariances will be zero, so their average will be zero. In other words, given the intuitively plausible idea that group selection is selection between groups and individual selection is selection between individuals within the *same* group, equation (2) tells us how much of the change in \bar{z} is due to group and individual selection, respectively.

The Contextual Approach

The second partitioning technique is motivated by an apparent shortcoming of the Price approach emphasized by Heisler and Damuth (1987); a similar point is made by Sober (1984) and Nunney (1985b). The shortcoming can be seen as follows. Consider again a population subdivided into groups of equal size. Suppose that the fitness of a given organism depends only on its own phenotypic character and not on the characters of others organisms in its group. Thus, any two organisms of identical phenotype have the same fitness, whichever groups they are found in—there are no

group effects on individual fitness. As Heisler and Damuth (1987) note, most biologists would say that no group selection is occurring in such a situation, for the evolution of the system can be predicted without taking group structure into account. Nonetheless, the term $\text{Cov}(Z_j, W_j)$ in equation (2) may very well be nonzero. Some groups may be fitter than others, simply because they contain a higher proportion of fitter organisms. So equation (2) detects a component of group selection where intuitively there is none.

As a solution to this problem, Heisler and Damuth (1987) advocate a statistical technique drawn from the social sciences called ‘contextual analysis’ (for a thorough introduction to contextual analysis see Boyd and Iversen 1979; for applications to levels of selection problems in biology, see Heisler and Damuth 1987; Goodnight et al. 1992; Tsuji 1995). The basic idea is to regard group character as a contextual or relational character of each individual in the group. Therefore, every individual is assigned two characters: an individual character z_{ij} and a contextual character Z_j . Both of these characters are capable of affecting an organism’s fitness. The crucial question vis-à-vis group selection is: Does the contextual character (group character) of an organism affect its fitness? Only if the answer is yes should we say that group selection is in operation—for only then is there a group-level effect on individual fitness.

To answer this question, we employ the following linear regression model:

$$w_{ij} = \beta_1 z_{ij} + \beta_2 Z_j + e_{ij}, \quad (3)$$

where β_1 is the partial regression of individual fitness on individual character, controlling for group character; β_2 is the partial regression of individual fitness on group character, controlling for individual character; and e_{ij} is the residual whose variance is to be minimized. Therefore, β_2 is the change in individual fitness that would result if the group character of an individual of fixed phenotype were changed by one unit—it measures the extent to which differences in group character predict differences in individual fitness, holding individual character constant.

Heisler and Damuth’s suggestion is that group selection requires that β_2 be nonzero. This means that information about the group to which an individual belongs is relevant to predicting the individual’s fitness, over and above information about the individual’s own phenotype, that is, it signals a group effect on individual fitness. In the situation where fitness is determined by individual phenotype alone, irrespective of group character, then clearly β_2 will equal zero, even though $\text{Cov}(Z_j, W_j)$ may be nonzero. So the contextual approach, unlike the Price approach, yields the intuitively correct result: there is no group selection in such a situation.

Like the Price approach, the contextual approach also partitions the total selection differential into two components, but differently:

$$\bar{w}\Delta\bar{z} = \beta_1 \text{Var}(z_{ij}) + \beta_2 \text{Cov}(z_{ij}, Z_j). \quad (4)$$

The first term on the right side of equation (4), $\beta_1 \text{Var}(z_{ij})$, is the partial regression of individual fitness on individual character times the variance in individual character; the second term, $\beta_2 \text{Cov}(z_{ij}, Z_j)$, is the partial regression of individual fitness on group character times the covariance between an

individual's character z_{ij} and the character of its group, Z_j . The two terms correspond respectively to the change due to individual selection and the change due to group selection, according to the contextual approach. Because equations (2) and (4) are different, it is obvious that the Price approach and the contextual approach conflict; they offer different estimates of the relative contributions of individual and group selection to total evolutionary change.

The relation between the Price and contextual approaches can be seen more clearly by rearranging equations (2) and (4). Since $Z_j = 1/n \sum_i z_{ij}$, it follows that $\text{Cov}(z_{ij}, Z_j) = \text{Var}(Z_j)$. Therefore, equation (4) can be written:

$$\bar{w}\Delta\bar{z} = \beta_1\text{Var}(z_{ij}) + \beta_2\text{Var}(Z_j). \tag{5}$$

Now recall equation (2), the Price partition, $\bar{w}\Delta\bar{z} = E_j[\text{Cov}_w(z_{ij}, w_{ij})] + \text{Cov}(Z_j, W_j)$. Because $W_j = 1/n \sum_i w_{ij}$, it follows that $\text{Cov}(Z_j, W_j)$ equals $\text{Cov}(Z_j, w_{ij})$, where the latter is the covariance between an individual's fitness and the group character of that individual. Second, in virtue of the general relation between covariance and regression, $\text{Cov}(Z_j, w_{ij}) = b_{wZ}\text{Var}(Z_j)$, where b_{wZ} is the simple regression of individual fitness on group character. Therefore, equation (2) can be written:

$$\bar{w}\Delta\bar{z} = E_j[\text{Cov}_w(z_{ij}, w_{ij})] + b_{wZ}\text{Var}(Z_j). \tag{6}$$

This form permits simple comparison between the Price and contextual approaches with respect to the criterion for group selection. On the Price approach, the condition for group selection is that $b_{wZ}\text{Var}(Z_j) \neq 0$. This, in turn, requires that $b_{wZ} \neq 0$. On the contextual approach, group selection requires $\beta_2 \neq 0$. Therefore, the Price approach says that group selection requires the simple regression of individual fitness on group character to be nonzero, whereas the contextual approach says that the partial regression of individual fitness on group character, controlling for individual character, must be nonzero.

Is the Price approach or the contextual approach superior? This is a tricky question. On the one hand, the contextual approach seems theoretically preferable. If we think of group selection and individual selection as separate evolutionary forces, each capable of influencing the evolution of a trait, then the use of partial regression techniques is surely appropriate. For the basic motivation behind the contextual approach—that an association between fitness and group character may be a mere side effect of the association of fitness and individual character—is difficult to deny; it is simply a special case of the familiar point that direct selection on a given character will have indirect effects on any correlated characters (Pearson 1903; Lande and Arnold 1983). Considering the partial regression of fitness on a character, controlling for correlated characters, rather than the simple regression, is the obvious way to accommodate this point. This is a point in favor of the contextual approach.

On the other hand, the contextual approach has the consequence that group selection can occur even if there is no variance in group fitness, as Goodnight et al. (1992) showed. These authors consider a model of soft selection, where each group is constrained to leave identical numbers of progeny, so the variance in group fitness is zero, but the fitness of any individual depends on its relative phenotypic ranking within

its group. The Price approach implies that there is no group selection in such a scenario, for $\text{Cov}(Z_j, W_j) = 0$, but the contextual approach detects a component of group selection, for β_2 is nonzero—there are group effects on individual fitness. Intuitively, this is an odd result. Surely, selection at any given hierarchical level requires variance in fitness at that level? This is a point in favor of the Price approach.

An extended comparison of the Price and contextual approaches to multilevel selection is attempted later in this paper. But first, I introduce a third partitioning technique, based on some ideas of Nunney (1985a,b).

The Neighbor Approach

In motivation, the neighbor approach is similar to the contextual approach; both agree that the Price approach is deficient because it allows group selection in the absence of group effects on individual fitness. As we saw, the contextual approach employs a multiple regression model to assess the effect of group character on fitness. The two independent variables in this model were individual character z_{ij} , and group character Z_j ; the regression model was $w_{ij} = \beta_1 z_{ij} + \beta_2 Z_j + e_{ij}$. The neighbor approach also uses a multiple regression model, but with a slight difference. The two independent variables in the neighbor model are individual character z_{ij} and neighborhood character h_{ij} , where an individual's neighborhood character is the average phenotypic character of its neighbors. An individual's neighbors are defined as all the organisms in its group except itself. The aim of the neighbor approach is to detect neighborhood effects on individual fitness, rather than group effects. The regression model is therefore:

$$w_{ij} = \beta_3 z_{ij} + \beta_4 h_{ij} + e_{ij}, \tag{7}$$

where β_3 is the partial regression of fitness on individual character, controlling for neighborhood character, and β_4 is the partial regression of fitness on neighborhood character, controlling for individual character.

There is an obvious relation between an individual's neighborhood character, h_{ij} , its individual character, z_{ij} , and its group character, Z_j . The relation is

$$h_{ij} = \frac{1}{n-1} \left[\left(\sum_i z_{ij} \right) - z_{ij} \right] = \frac{1}{n-1} [nZ_j - z_{ij}]. \tag{8}$$

This means that there is a simple relation between β_4 , the partial regression of fitness on neighborhood character, and β_2 , the partial regression of fitness on group character. β_2 measures the increase in fitness when the group character of an individual increases by one unit, its individual character remaining the same; while β_4 measures the increase in fitness when the neighborhood character of an individual increases by one unit, its individual character remaining the same. But when an individual's group character increases by one unit, its individual character remaining the same, then its neighborhood character must increase by $n/(n-1)$ units, because $\delta h/\delta Z = n/(n-1)$. Therefore,

$$\beta_2 = n/(n-1)\beta_4. \tag{9}$$

Because $n > 0$, this means that β_2 and β_4 always have the same sign. This is highly intuitive: the only way that there

can be group effects on individual fitness is if there are neighborhood effects on individual fitness and vice versa.

Because β_2 and β_4 are so closely related, and the difference between them will be very small if n is large, it might be wondered whether there is really any substantive difference between the contextual and neighbor approaches. Surprisingly, there is. The difference becomes apparent when we look, not at the regression coefficients, but at how the neighbor approach partitions the total selection differential on z :

$$\bar{w}\Delta\bar{z} = \beta_3\text{Var}(z_{ij}) + \beta_4\text{Cov}(z_{ij}, h_{ij}). \quad (10)$$

According to the neighbor approach, the two terms on the right side of equation (10) correspond respectively to individual and group selection. The key difference behind the neighbor approach and the contextual approach lies in the second term, which reflects group selection. On the neighbor approach, this term is $\beta_4\text{Cov}(z_{ij}, h_{ij})$; on the contextual approach, the analogous term is $\beta_2\text{Cov}(z_{ij}, Z_j)$. As we saw, $\text{Cov}(z_{ij}, Z_j) = \text{Var}(Z_j)$, so is always positive unless the groups all have the same group character, in which case it is zero. That is, there is an intrinsic correlation between individual character and group character. But the term $\text{Cov}(z_{ij}, h_{ij})$ is very different. This term is the covariance between individual character and neighborhood character, and it may take on any value. If groups are formed at random, then $\text{Cov}(z_{ij}, h_{ij})$ will equal zero—there will be no correlation between an individual's own phenotypic character, z_{ij} , and the average phenotypic character of its neighbors, h_{ij} . If groups are formed assortatively, that is, if individuals with similar phenotypic characters tend to be found in the same groups, then $\text{Cov}(z_{ij}, h_{ij})$ will be positive. So $\text{Cov}(z_{ij}, h_{ij})$ measures departures from randomness in the formation of groups; equivalently, $\text{Cov}(z_{ij}, h_{ij})$ measures the amount of phenotypic correlation between group members.

This means that, despite the close relation between the regression coefficients β_2 and β_4 , the contextual and neighbor approaches embody very different conceptions of group selection. The neighbor approach implies that group selection can only operate if there is above-random or below-random phenotypic variation between groups (a conclusion explicitly endorsed by Nunney 1985a,b). The contextual approach contains no such implication.

Is the contextual approach or the neighbor approach superior? Again, this is a tricky question. Some authors have explicitly rejected the idea that group selection requires non-random formation of groups. Thus, for example, Sober and Wilson (2002) criticized some multilevel selection theorists for “focusing exclusively on correlation, which erroneously gives the impression that group selection requires above-random variation among groups” (p. 720). By contrast, Nunney (1985b) explicitly said that a trait can only spread or be maintained by group selection if there is “positive association of individuals exhibiting the trait” (p. 221). So, there is no consensus on this point in the literature.

One possible reason for favoring the neighbor over the contextual approach stems from a general consideration about the likely causal influences on individual fitness. Statistical techniques such as multiple regression do not in themselves tell us what causal factors are influencing the response variable; correlation is not causation. So, in general, successful

use of regression analysis requires independent causal knowledge. It could perhaps be argued that the neighborhood character of a given individual, h_{ij} , is the sort of character that could directly causally influence the individual's fitness, but that the group character, Z_j , is not. Put differently, any causal link between individual fitness and group character must necessarily be indirect, mediated by a direct causal link between individual fitness and neighborhood character. If an individual's fitness is affected by its interactions with its neighbors, that is, its fellow group members, then there will be a correlation between individual fitness and neighborhood character and hence between individual fitness and group character. But arguably, only the former correlation is causal. Therefore, neighborhood character, rather than group character, is the appropriate independent variable in the multiple regression model. Thus, the neighbor approach to multilevel selection is better than the contextual approach.

This is not intended as a knock-down argument, only a consideration that might help choose between the contextual and neighbor approaches. As noted above, the contextual and neighbor approaches are very similar in spirit—both are motivated by the same shortcoming of the Price approach—and in practice the difference between the two will often be negligible. For it is easy to show that:

$$\beta_4\text{Cov}(z_{ij}, h_{ij}) = \beta_2\text{Var}(Z_j) - \frac{1}{n}\beta_2\text{Var}(z_{ij}). \quad (11)$$

As group size n increases, the difference between $\beta_4\text{Cov}(z_{ij}, h_{ij})$ and $\beta_2\text{Var}(Z_j)$ will therefore tend to zero. In practice, the contextual and neighbor approaches will usually partition $\Delta\bar{z}$ into very similar, though not identical, components. Nonetheless, form a theoretical point of view the contextual and neighbor approaches do differ in an important way.

One final point of difference between the contextual and neighbor approaches is worth noting: the former is more generally applicable than the latter. In outlining the contextual approach above, we defined group character Z_j to be the mean character of the individuals in the j th group, $Z_j = 1/n \sum_i z_{ij}$. However, the contextual approach is applicable even if group character is not defined this way, but is instead an emergent character of the group (e.g., population density) that is not the mean value of any character measurable on individuals. By contrast, the neighbor approach is not applicable in such circumstances. (It is not intended to be.) By definition, an individual's neighborhood character h_{ij} is the mean phenotypic character of its neighbors; so the choice between using group character Z_j or neighborhood character h_{ij} in the regression model only arises where $Z_j = 1/n \sum_i z_{ij}$, that is, where Z_j is not an emergent group character. Hence, the contextual approach has greater generality. In what follows, it is assumed that we are not dealing with emergent group characters, hence the choice between contextual and neighbor approaches does arise.

The Three Approaches Compared

So there are three different approaches to partitioning the total change in \bar{z} into components that, allegedly, correspond to individual and group selection. To facilitate comparison, the three partitions are:

	Individual selection	Group selection
Price approach (see eq. 2)	$\bar{w}\Delta\bar{z} = E_j[\text{Cov}_w(z_{ij}, w_{ij})] +$	$\text{Cov}(Z_j, W_j),$
contextual approach (see eq. 5)	$\bar{w}\Delta\bar{z} = \beta_1\text{Var}(z_{ij}) +$	$\beta_2\text{Var}(Z_j),$
and		
neighbor approach (see eq. 10)	$\bar{w}\Delta\bar{z} = \beta_3\text{Var}(z_{ij}) +$	$\beta_4\text{Cov}(z_{ij}, h_{ij}).$

How should we choose between the three approaches? One might regard this question as semantic, a matter of deciding how the terms “individual selection,” “group selection” and “multilevel selection” should be used. But the issue is not this trivial. The concept of group selection, in particular, has a long and controversial history in evolutionary discussions, from Darwin to the present day. So it is important to ask which of the three partitioning techniques most accurately captures the concept that the participants in these discussions have had in mind. Therefore, the choice between equations (2), (5) and (10) does not boil down to a matter of definition. The traditional explanatory role of the concepts of group and individual selection imposes constraints on what an adequate partitioning technique can look like.

One way to compare the three approaches is to contrast the answers that each approach gives to a range of critical questions. Five such questions, and the respective answers to them, are given below.

(1) *Does group selection require variance in group fitness?*—The Price approach answers yes: without variance in group fitness, $\text{Cov}(Z_j, W_j)$ will obviously be zero, as noted above. However, the other two approaches say no. As we saw, Goodnight et al. (1992) showed that in soft-selection models with no variance in group fitness, the partial regression coefficient β_2 is nonzero, hence there is a component of group selection according to the contextual approach. Because β_2 and β_4 always have the same sign, the neighbor approach will also detect a component of group selection in such models so long as the groups are formed nonrandomly.

(2) *Does group selection require variance in group character?*—The Price approach says yes: If all groups have the same group character, then $\text{Cov}(Z_j, W_j)$ must be zero. The contextual approach agrees: If all groups have the same character, then $\text{Var}(Z_j)$ will be zero, hence $\beta_2\text{Var}(Z_j)$ will be zero, too. Equivalently, β_2 , the partial regression of fitness on individual character holding group character fixed, will equal the simple regression of fitness on individual character, because all groups have the same character. Note that this is the only circumstance in which the Price and contextual approaches agree about the amount of evolutionary change attributable to group selection.

On the neighbor approach, matters are somewhat different. If there is no variance in group character, then the partial regression coefficients β_3 and β_4 are not well defined, so the partitioning technique cannot be applied. This is because zero variance in group character means that individual character z_{ij} and neighborhood character h_{ij} must be perfectly correlated. (Recall that $h_{ij} = (1/n - 1)[nZ_j - z_{ij}]$; so if Z_j is constant for all j , then h_{ij} is a linear function of z_{ij} , hence h and z are correlated perfectly.) The point here is a technical one: Where

two independent variables in a multiple regression analysis are perfectly correlated, the partial regression coefficients are not well defined. (This is known as the problem of perfect collinearity in statistics.) Intuitively, this is easy to understand. If z and h correlate perfectly, then we cannot ask about the effect on fitness of changing z while keeping h fixed; it is impossible to change z while keeping h fixed. So β_3 is undefined; similarly for β_4 . Therefore, in this situation, the neighborhood approach cannot be applied—the total selection differential cannot be partitioned into two components.

(3) *Does group selection require nonrandom formation of groups?*—Both the Price and contextual approaches say no. As long as the groups vary in character, then there can be group selection, according to both. Random assortment of organisms into groups generates variance in group character, just as random assortment of alleles into diploid genotypes (i.e., random mating) generates genotypic variance. But the neighbor approach says yes, as we saw in the previous section. If groups are formed at random, there will be no correlation between an individual’s character and that of its neighbors, so $\text{Cov}(z_{ij}, h_{ij})$ will be zero. (Even with random group formation, there will of course be a correlation between an individual’s character and that of its group, so $\text{Cov}(z_{ij}, Z_j) = \text{Var}(Z_j)$ will be nonzero.) This is the distinctive feature of the neighbor approach.

(4) *Does group selection require individual fitnesses to be group-dependent?*—The Price approach says no, as we have seen. Even if an individual’s fitness is a function of its own phenotype alone and is wholly unaffected by the composition of its group, the term $\text{Cov}(Z_j, W_j)$ can be nonzero. It is precisely to avoid this result that the contextual and neighbor approaches insist on looking at the partial regression of fitness on group character/neighborhood character. Because both β_2 and β_4 are zero in the situation where there are no group effects on individual fitness, the contextual and neighbor approaches answer no to this question.

(5) *Does individual selection require within-group variance in fitness?*—The Price approach says yes: If within each group, all individuals have the same fitness, then each of the $\text{Cov}_w(z_{ij}, w_{ij})$ terms will be zero, so their average, $E_j[\text{Cov}_w(z_{ij}, w_{ij})]$, will be zero, too. The contextual approach agrees. In the absence of any within-group variance in individual fitness, the term β_1 will also be zero. This is because β_1 measures the fitness difference between an individual of fixed group character when that individual’s own character is increased by one unit, that is, the fitness difference between two individuals within the same group whose individual characters differ by one unit. But if within-group fitnesses are identical, then this difference must equal zero. However, the neighbor approach disagrees, for the term β_3 can be nonzero even if there is no within-group variance in fitness. This is because β_3 measures the fitness difference between an individual of fixed neighborhood character when its individual character is increased by one unit. But two individuals of identical neighborhood character whose individual characters differ cannot be in the same group, so the fact that there is no within-group variance does not have any effect on β_3 .

A simple example may help make this last point clear. Suppose that individuals are of two types, A and B, and groups are of size $n = 4$. We define z_{ij} as one if the i th

organism in the j th group is A, zero otherwise; so \bar{z} is the overall frequency of type A. Suppose there are only two types of groups in the population, AAAB and AAB, whose relative frequencies are equal. The fitness of each type of individual (w_{ij}) in each type of group is 5, 5, 5, 5 for AAAB and 6, 6, 6, 6, for AAB. Note that within any group, all individuals have the same fitness.

To see the difference between the contextual and neighbor approaches, pick one of the A individuals in an AAAB group. This individual has $w_{ij} = 5$, $z_{ij} = 1$, $Z_j = 3/4$, $h_{ij} = 2/3$. The coefficient β_1 measures the fitness difference between this individual and the fitness of a B individual with the same group character Z , that is, in the same group. This difference is obviously zero—the A and B individuals in an AAAB group have the same fitness, by design. But the coefficient β_3 measures the fitness difference between our focal A individual and a B individual with the same neighborhood character h , that is, a B individual in an AAB group. For an A in an AAAB group and a B in an AAB group have the same neighborhood character—both have AAB neighbors. So $\beta_3 = 6 - 5 = 1$. In other words, the contextual approach detects individual selection by comparing the fitnesses of individuals within groups, for such individuals have the same group character, while the neighbor approach compares the fitnesses of individuals across groups, for such individuals have the same neighborhood character.

Group Selection: A Concept in Tension

Which of the three approaches gives the most plausible answers to our five questions? With respect to question 1, the Price approach seems clearly superior to the other two. The idea that group selection can operate even if all groups have identical fitness is intuitively very strange, for it violates the famous dictum that selection at any given hierarchical level requires heritable variance in fitness at that level (Lewontin 1970). (More on this point in a moment.) On the other hand, the Price approach's answer to question 4 seems wrong: if an individual's fitness is a function of its own phenotype alone, and does not depend on its group, then most biologists would insist that no group selection can occur. This is a point in favor of the contextual and neighbor approaches.

With respect to question 3, honors are divided. The importance of nonrandom assortment for the evolution of altruism has long been recognized; and altruism and group selection are closely linked subjects. So, a conception of group selection that requires nonrandom formation of groups, i.e., correlations between group members, fits reasonably well with the traditional explanatory role of the concept. On the other hand, if we allow that the genotypic variance generated by random mating is sufficient for individual (genotypic) selection to take place, it seems somewhat odd to insist that nonrandom formation of groups is required for group selection. For as many authors have noted, there is a clear analogy between the fusion of two genes to form a diploid genotype and the coming together of individuals to form groups. Indeed, the formal apparatus of multilevel selection theory applies neatly to diploid genetics, simply by thinking of a diploid organism as a group of two genes (Wilson 1990; Kerr

and Godfrey-Smith 2002). Thus, there is something to be said in favor of both answers to question 3.

With respect to question 5, the Price and contextual approaches score better than the neighbor approach. The idea that individual selection can operate in the absence of within-group variance in fitness is certainly odd, just as the idea that group selection can operate in the absence of between-group variance in fitness is odd. With respect to question 2, there is little to choose between the three approaches, since the Price and contextual approaches answer identically, and intuitively correctly, while the neighbor approach fails to deliver an answer, as the relevant regression coefficients are not defined.

I suggest that the 'best' answers to our five questions, in the sense of corresponding most closely to the concepts that evolutionists have in mind when they talk about individual and group selection, are probably: question 1, yes; question 2, yes; question 3, unclear; question 4, yes; and question 5, yes.

None of three partitioning techniques generates this set of answers, even given the latitude that the unclear status of question 3 permits. The crucial trade-off appears to be between questions 1 and 4. The whole motivation behind the contextual and neighbor approaches is that the Price approach generates the intuitively incorrect answer to question 4—it says that there can be group selection even if there are no group effects on individual fitness. The contextual and neighbor approaches are designed to, and do, generate the correct answer to question 4, but at a cost. For while the Price approach gets the answer to question 1 right, the contextual and neighbor approaches do not—both have the consequence that group selection can operate even if all groups have identical fitness.

One might argue that this consequence—group selection without variance in group fitness—is not so strange after all, by invoking the well-known distinction between the two different types of multilevel selection (Arnold and Frisrup 1982; Damuth and Heisler 1988). Following Damuth and Heisler's terminology, in multiselection level 1 (MLS1) group fitness is defined as average individual fitness, as in the three models above, while in multilevel selection 2 (MLS2) group fitness is defined differently, as expected number of offspring groups. MLS1 and MLS2 are fundamentally different in intent: in the former, individual organisms are the entities of interest, and groups are relevant only in so far as they affect the evolution of individual phenotypes; in the latter, the groups themselves are the entities of interest. (As Damuth and Heisler [1988] note, virtually all the work on group selection by geneticists is concerned with MLS1; the macroevolutionary concept of species selection is the best-known example of MLS2.) It might be argued that the Lewontin (1970) criteria apply only to MLS2 and not to MLS1, that is, that when Lewontin says that selection at any given hierarchical level requires heritable variance in fitness at that level, the relevant sense of fitness is the MLS2 sense. If this were correct, then the fact that the contextual and neighbor approaches permit group selection without variance in group fitness would not count against them, for the sense of fitness in which they permit this is the MLS1 sense, that is, average individual fitness.

This argument seems right—and it is certainly true that MLS2 requires the Lewontin criteria, in the MLS2 sense of fitness—but it cannot be the whole story. For a number of theorists whose concern is MLS1 and who are clearly aware of the distinction between MLS1 and MLS2, nonetheless maintain that group selection requires variance in group fitness in the MLS1 sense (average individual fitness). For example, Sober and Wilson (1998) offer a general account of how altruism can evolve by group selection, taking the Price equation as their conceptual tool; they explicitly define group fitness as average individual fitness and argue that group selection can only occur if there is variance in group fitness (p. 113). Similarly, Wade et al. (1999), in a discussion of interdemic (group) selection, acknowledge that “most would consider [variance in demic productivity] a prerequisite for interdemic selection” (p. 600); but in pure soft selection, contextual analysis detects interdemic selection even though all demes have identical productivity (fitness), as they note. So the apparently unpalatable consequence of the contextual and neighbor approaches, that they permit group selection without variance in group fitness, cannot be defused simply by appealing to the distinction between MLS1 and MLS2. Even in studies where the focus is clearly on MLS1 not MLS2, variance in group fitness is often regarded as necessary for group selection.

This suggests that the traditional concept of group selection constitutes a shotgun marriage between two separate ideas, and so is inherently unstable. The joint requirements that group selection involve variance in group fitness (in the MLS1 sense) and that it involve group effects on individual fitness are in serious tension with one another. The former requirement reflects the idea that group selection should be analogous to ordinary individual selection, which obviously requires variance in individual fitness; by parity of argument, surely group selection requires variance in group fitness. The latter requirement reflects the idea that higher-level units of selection emerge from fitness-affecting interactions between lower-level units. The close association, historical and conceptual, between group selection and altruism reinforces this point. Where altruistic behaviors are in question, there are by definition group effects on individual fitness. The intuition that real group selection cannot operate in the absence of group effects on individual fitness, even if $\text{Cov}(Z_j, W_j)$ is nonzero, probably stems from this traditional association between altruism and group selection in evolutionary discussions.

The tension between the variance in group fitness requirement and the group effects on individual fitness requirement is reflected by the fact that none of the three partitioning techniques generates the correct answers to all five of our critical questions. The Price approach respects the former requirement, but not the latter; the contextual and neighbor approaches respect the latter requirement, but not the former. If we insist that group selection satisfy both requirements, we seem forced to conclude that there is no way of partitioning the overall selection differential on a character into components corresponding to individual and group selection. This is a profoundly unsatisfactory conclusion. For, in the absence of such a partitioning technique, it makes little sense to ask whether individual or group selection is a more pow-

erful force in a given situation; and evolutionists have long assumed that that question does make sense.

One possible way of avoiding this conclusion and reconciling the requirements of variance in group fitness and group effects on individual fitness, is to impose constraints on what counts as a group. None of the three partitioning techniques involves a commitment to any particular definition of a group; indeed, each would work perfectly well even if the groups were arbitrarily defined by the investigator and had no biological reality at all. (The question of how to define groups is independent of the question of which of the three partitioning techniques is preferable.) So something needs to be said about what groups are. D. S. Wilson has argued persuasively that groups should be defined on the basis of fitness interactions—an organism’s group consists of all those other organisms with which it engages in fitness-affecting interactions (Wilson 1975; Sober and Wilson 1998); Uyenoyama and Feldman (1980) have a similar conception. Following Wilson, I call this the trait-group conception of a group.

If we adopt the Price approach to multilevel selection in conjunction with the trait-group conception of a group, then we generate the correct answer to questions 1 and 4 (Okasha 2004). (The Price approach automatically gets the answer to question 1 correct; question 4 is the stumbling block.) For in a situation where there are no group effects on individual fitness, that is, where an individual’s fitness is a function of its own phenotype alone, then it follows that there are no trait groups in the population, and thus that there is no possibility of group selection. For trait groups by definition require fitness-affecting interactions between group members, they require that individual fitness is affected by group membership. So the problem with the Price approach that the contextual and neighbor approaches were designed to remedy is solved by extrastatistical means; the result is a conception of multilevel selection that gets the answers to questions 1 and 4 right. Group selection requires both variance in group fitness, for otherwise $\text{Cov}(Z_j, W_j) = 0$, and group effects on individual fitness, for otherwise there are no groups in the population, hence obviously no possibility of group selection.

This move may seem attractive, but in fact it conceals, rather than resolves, the tension between the requirements of variance in group fitness and group effects on individual fitness. For the theoretical argument in favor of the contextual and neighbor approaches is not just that they yield the correct answer in the limiting case where there are no group effects on individual fitness. The failure of the Price approach to get the answer right in this case is symptomatic of a deeper problem, namely that the term $\text{Cov}(Z_j, W_j)$ will always contain an element that reflects the fact that individuals whose own phenotypes make them fitter than average will tend to be found in the fitter groups. This remains true even if there are group effects on individual fitness. So, although employing the Price approach in conjunction with the trait-group definition of group avoids the problem in the limiting case, it remains the case that the term $\text{Cov}(Z_j, W_j)$, which proponents of the Price approach attribute to group selection, assumes the value it does in part because of the intrinsic correlation between individual character and group character. The use of partial regression techniques is the natural way to correct for this correlation. Thus, the theoretical argument against

the Price approach, and in favor of the contextual or neighbor approaches remains intact, even if we employ the Price approach in conjunction with the trait-group concept.

This point can be seen from another angle. Suppose that trait-groups do exist in a population, but that the fitness-affecting interactions within them are extremely weak. An organism's fitness depends almost entirely on its own phenotypic character and only very slightly on the character of its group. Intuitively, the amount of change attributable to group selection in this situation should be very small, for the situation differs only marginally from the case where there are no group effects on individual fitness at all. The contextual and neighbor approaches will reveal a very small group selection component, for β_2 and β_4 will be close to zero. But it is quite possible that $\text{Cov}(Z_j, W_j)$ will be very high. If there is little fitness variance within groups, then $\text{Cov}(Z_j, W_j)$ will be a very high fraction of the total change $\Delta\bar{z}$, so the Price approach will attribute most of the change to group selection. Combining the Price approach with the trait-group definition of group avoids the consequence that group selection can operate where there are no group effects on individual fitness, but it still allows that where there are very weak group effects on individual fitness, group selection may be extremely strong. Therefore, the theoretical argument in favor of the contextual and neighbor remains: $\text{Cov}(Z_j, W_j)$ contains a component that is not due to group selection, intuitively.

Conclusion

The foregoing analysis has two main consequences. First, it shows that the concept of multilevel selection is less clear-cut than is sometimes assumed; second, that the traditional concept of group selection involves two requirements that are in serious tension with one another. The idea that the evolution of a trait, or a gene, can be influenced by selection at more than one level of the biological hierarchy is a familiar one; so also is the idea that it is possible to compare the magnitude of the selective forces at different levels. This leads naturally to the idea of partitioning the overall selection differential on a character into separate components, each corresponding to a different level of selection. But there are three nonequivalent partitioning techniques and no obvious way of choosing between them, suggesting that the concept of multilevel selection is ambiguous. (This ambiguity is additional to the ambiguity stemming from the MLS1/MLS2 distinction; it arises from within an MLS1 framework.)

Where the levels of selection we are dealing with are individuals and groups, we might hope to decide between the three partitions by appealing to the traditional concept of group selection. But that concept involves two distinct aspects—variance in group fitness and group effects on individual fitness. None of the three partitioning techniques captures both aspects. This suggests that the traditional concept of group selection involves a tension. There simply may not be any well-defined evolutionary process affecting natural populations that satisfies both the requirements of variance in group fitness and group effects on individual fitness. If by a well-defined process, we mean a process that has a quantifiable effect on the total evolutionary change, that is the conclusion we seem driven inescapably toward.

ACKNOWLEDGMENTS

For comments and discussion of this material, I thank D. Sloan Wilson, E. Sober, P. Godfrey-Smith, B. Kerr, M. Wade, L. Lloyd, M. Fagan, an anonymous reviewer, and especially J. Damuth. I thank the Arts and Humanities Research Board (U.K.) for financial support.

LITERATURE CITED

- Arnold, A. J., and K. Fristrup. 1982. The theory of evolution by natural selection: a hierarchical expansion. *Paleobiology* 8: 113–129.
- Boyd, L. H., and G. R. Iversen. 1979. Contextual analysis: concepts and statistical techniques. Wadsworth, Belmont, CA.
- Frank, S. 1999. *Foundations of social evolution*. Princeton Univ. Press, Princeton, NJ.
- Goodnight, C. J., J. M. Schwartz, and L. Stevens. 1992. Contextual analysis of models of group selection, soft selection, hard selection, and the evolution of altruism. *Am. Nat.* 140:743–761.
- Hamilton, W. D. 1975. Innate social aptitudes in man: an approach from evolutionary genetics. Pp. 133–155. *in* R. Fox, ed. *Bio-social anthropology*. Wiley, New Jersey.
- Heisler, I. L., and J. Damuth. 1987. A method for analyzing selection in hierarchically structured populations. *Am. Nat.* 130: 582–602.
- Kerr, B., and P. Godfrey-Smith. 2002. Individualist and multi-level perspectives on selection in structured populations. *Biol. Philos.* 17:477–517.
- Lande, R., and S. J. Arnold. 1983. The measurement of selection on correlated characters. *Evolution* 37:1210–1226.
- Lewontin, R. 1970. The units of selection. *Annu. Rev. Ecol. Syst.* 1:1–18.
- Maynard Smith, J., and E. Szathmary. 1995. *The major transitions in evolution*. Oxford Univ. Press, Oxford, U.K.
- Michod, R. 1997. Evolution of the individual. *Amer. Natur.* 150: S5–S21.
- . 1999. *Darwinian dynamics*. Princeton Univ. Press, Princeton, NJ.
- Nunney, L. 1985a. Female-biased sex ratios: individual or group selection? *Evolution* 39:349–361.
- . 1985b. Group selection, altruism, and structured-deme models. *Am. Nat.* 126:212–230.
- Okasha, S. 2004. Multi-level selection, covariance and contextual analysis. *Br. J. Philos. Sci.* *In press*.
- Pearson, K. 1903. *Mathematical contributions to the theory of evolution*. XI. On the influence of natural selection on the variability and correlation of organs. *Philos. Trans. R. Soc. Lond.* A200: 1–66.
- Pomiankowski, A. 1999. Intra-genomic conflict. Pp. 121–152 *in* L. Keller, ed. *Levels of Selection in Evolution*. Princeton Univ. Press, Princeton, NJ.
- Price, G. 1972. Extension of covariance selection mathematics. *Ann. Hum. Genet.* 35:485–490.
- Sober, E. 1984. *The nature of selection*. Univ. of Chicago Press, Chicago.
- Sober, E., and D. S. Wilson. 1998. *Unto others: the evolution and psychology of unselfish behaviour*. Harvard Univ. Press, Cambridge, MA.
- . 2002. Reply to commentaries. *Philos. Phenomenol. Res.* 67(3):711–727.
- Tsuji, K. 1995. Reproductive conflicts in the ant *Pristomyrmex pungens*: contextual analysis and partitioning of covariance. *Am. Nat.* 154:599–613.
- Uyenoyama, M., and M. W. Feldman. 1980. Theories of kin and group selection: a population genetics perspective. *Theoret. Popul. Biol.* 17:380–414.
- Wade, M. 1985. Soft selection, hard selection, kin selection, and group selection. *Am. Nat.* 125:61–73.
- Wade, M., C. J. Goodnight, and L. Stevens. 1999. Design and interpretation of experimental studies of interdemec selection: a reply to Getty. *Am. Nat.* 154:599–603.

- Wilson, D. S. 1975. A theory of group selection. *Proc. Natl. Acad. Sci. USA* 72:143–146.
- . 1990. Weak altruism, strong group selection. *Oikos* 59: 135–140.
- . 1997. Altruism and organism: disentangling the themes of multi-level selection theory. *Am. Nat.* 150:S122–S134.

Corresponding Editor: S. Gavrilets