Multi-level selection and the issue of environmental homogeneity

Ciprian Jeler

(Department of Interdisciplinary Research – Humanities and Social Sciences,
"Alexandru Ioan Cuza" University of Iaşi, Romania)

Abstract

In this paper, I identify two general positions with respect to the relationship between environment and natural selection. These positions consist in claiming that selective claims need and, respectively, need not be relativized to homogenous environments. I then show that adopting one or the other position makes a difference with respect to the way in which the effects of selection are to be measured in certain cases in which the focal population is distributed over heterogeneous environments. Moreover, I show that these two positions lead to two different interpretations – the Pricean and contextualist ones – of a type of selection scenarios in which multiple groups varying in properties affect the change in the metapopulation mean of individual-level traits. Showing that these two interpretations stem from different attitudes towards environmental homogeneity allows me to argue: a) that, unlike the Pricean interpretation, the contextualist interpretation can only claim that drift or selection is responsible for the change in frequency of the focal trait in a given metapopulation *if* details about whether or not group formation is random are specified; b) that the traditional main objection against the Pricean interpretation – consisting in arguing that the latter takes certain side-effects of individual selection to be effects of group selection – is unconvincing. This leads me to suggest that the ongoing debate about which of the two interpretations is preferable should concentrate on different issues than previously thought.

Keywords: philosophy of biology, natural selection, environmental homogeneity, multi-level selection, group selection, Robert Brandon

1. Introduction

In this paper, I distinguish between two general positions with respect to the relationship between environment and natural selection. These positions consist in adopting and, respectively, rejecting the idea that selective claims need to be relativized to homogenous environments. But this is not just a theoretical distinction with no concrete consequences. On the contrary, I will show that adopting one or the other position affects the way in which one should measure the effects of selection in certain cases in which the focal population is distributed over heterogeneous environments.

Moreover, I will argue that this distinction can cast a new light on the much disputed issue of group selection, and in particular on a type of multi-level selection case that has received much attention in recent decades because it seems to provide clues about the evolution of altruistic traits. This type of case – called MLS1 (Damuth and Heisler 1988; Okasha 2006) – involves multiple groups, concerns properties of groups that are obtained by averaging the values of particular traits possessed by the individual members of those groups, and the target of interest is the change in the average value of these individual-level traits in the metapopulation. But there are two possible ways of understanding MLS1, and I will call them – for reasons that will be indicated shortly – "the Pricean interpretation" (based on Price 1972) and "the contextualist interpretation" (based on Heisler and Damuth 1987). Philosophers have recently debated about which of the two interpretations of MLS1 is preferable (Okasha 2004a, b, 2006, 2011; Sober 2011, 2015; Earnshaw 2015; Jeler 2016).[1] On the biological side, Frank (2012) has attempted to defuse this debate, but I will show below why I do not think he has succeeded. For now, note that, despite Frank's argument, a recent discussion in the *Journal of Evolutionary Biology* (Gardner 2015a, b; Goodnight 2015) revolves around the same issue of deciding which of the two interpretations is more appropriate for MLS1 scenarios.

Roughly speaking, the main difference between the two competing interpretations is usually presented as follows: the Pricean interpretation understands group selection as the differential growth of groups (with some groups growing larger than others), whereas the contextualist interpretation sees it as the differential effect of group membership on individual fitnesses (e.g. Okasha 2006; Earnshaw 2015; McLoone 2015; Sober 2015[2]). I will argue below that this presentation of the two interpretations is incomplete. More specifically, I show that one of the main and, as yet, unacknowledged differences between the two interpretations consists in the fact that the Pricean interpretation adopts the idea that selective claims need to be relativized to homogenous environments, whereas the contextualist interpretation rejects this idea. Because I render explicit hitherto unacknowledged aspects of these interpretations, I prefer to also replace the terms generally used in the literature (i.e. the "Price" and "contextual approaches") with the "Pricean" and "contextualist" *interpretations*; this terminological replacement also helps preclude potential confusions.

---

[1] Since this debate only concerns MLS1 cases, this paper will do so as well, leaving aside the type of cases known as MLS2 (in which the target of interest is the change in the average value of group-level traits). Moreover, since they are not integral to the focal debate, this paper also leaves aside MLS1 cases that involve group characters like population density, division of labor etc., which are not mere averages of the individual traits within the group.

[2] Sober (2011) seemed somewhat reluctant to agree that the two approaches define group selection differently; but that no longer seems to be the case, as he recently stated (Sober 2015, 836) that "the two approaches disagree about what group selection *is*", before going on to note that his own "definitions of group and individual selection are compatible with the Price approach but not with the contextual approach".

Seeing the two interpretations in this new light brings both good and bad news for supporters of group selection in MLS1 scenarios. The bad news, as I will argue in Section 5, is that, according to the contextualist interpretation, groups are not units on which selection acts in MLS1 scenarios: what is called "group selection" in the contextualist interpretation is nothing more than selection on an individual-level trait (and it is therefore nothing else than individual selection); moreover, I will show that it may only be called "selection" *if* the uneven distribution of individual types into groups is not accidental, but is due to a trait of types (e.g. assortative preferences).

Therefore, a supporter of group selection in MLS1 scenarios is only left with the solution of embracing the Pricean interpretation. The good news though is that once we bring the issue of environmental homogeneity into this discussion, the main objection to the Pricean interpretation (that many authors have, for more than three decades now, considered to be devastating for the Pricean interpretation) becomes, in fact, unconvincing (Section 6). Once this objection – which consists in arguing that the Pricean interpretation takes certain side-effects of individual selection to be effects of group selection – is shown to be unconvincing, I will argue that the debate about which of the two interpretations is preferable (and, implicitly, about whether or not groups may be seen as units undergoing selection in MLS1 cases) should shift its focus towards assessing the validity of the two conjoint assumptions of the Pricean interpretation: namely the idea that selective claims should be relativized to homogenous environments and the idea that we can legitimately define a group's fitness as the mean fitness of its individual members (Section 6).

Before getting to these points, in Sections 2 and 3 I will detail the distinction between adopting and rejecting the idea that selective claims should be relativized to homogenous environments. And in Section 4, I provide a detailed MLS1 example that will be used in order to substantiate the claims laid out in the subsequent sections.

2. Selection and environments

Throughout this paper, I will rely on two central notions: the "Causal Condition" embedded in the definition of natural selection and the "environmental homogeneity restriction" of this condition. Section 2.1 will present the former, while Sections 2.2 and 2.3 will explain what I mean by the second notion and how this restriction can be instantiated by using Robert Brandon's notion of homogenous (selective) environments. Section 3 will then highlight the effects of adopting and rejecting this restriction in a case involving single-level selection.

2.1. The Causal Condition embedded in the definition of natural selection

Consider the difference between the following (admittedly incomplete) attempts at defining natural selection:

(A)   Natural selection is the process that occurs when the differential fitness of the focal biological entities is associated with their phenotypic differences;

(A*)   Natural selection is the process that occurs when the differential fitness of the focal biological entities is due to their phenotypic differences.

It is important to bear in mind that the "differential fitness" referred to here is *actual* or *realized* fitness. Now, taken as a definition, statement (A) is obviously unsatisfactory insofar as it fails to distinguish natural selection from other factors of evolutionary change such as drift, migration etc. Indeed, any factor that might be responsible for an association between variation in fitness and variation in phenotype (chance events or migration processes included) is relevant for statement (A). In contrast, if we want a definition of natural selection that does distinguish it from drift, migration and the like, we need to resort to a statement like (A*). We are dealing with natural selection only when variation in realized fitness and phenotype are not merely associated, but causally linked (in the sense that variation in phenotype causes variation in realized fitness).[3] I call this the *Causal Condition* embedded in the definition of natural selection.[4] The idea here is not to provide a complete definition of natural selection, but to point out that a definition of natural selection that pretends to be complete needs to (explicitly or implicitly) include this Causal Condition.

The Causal Condition does not commit us to any particular account of causation and it can be formulated by using a variety of existing accounts. Since I adapted his term of "Causal Condition" for my goals here, it is convenient to instantiate this notion by also using Glymour's (2014, 107) definition of "direct causal relations". Here it is: "A variable P is said to be a direct cause of a variable W relative to some set of variables ($\mathbf{V}$, P, $W \in \mathbf{V}$) and background conditions $\mathbf{B}$ when there is some pair of interventions on P, holding all other variables except W in $\mathbf{V}$ constant, such that the probability distribution or density over W differs across the interventions." By replacing variable P with "a difference in a particular phenotypic trait" and variable W with "a difference in fitness", we get the Causal Condition.

2.2. The "environmental homogeneity restriction" of the Causal Condition

In biological theory and practice, a number of restrictions to the Causal Condition seem to be routinely imposed. Some set-ups might (arguably) satisfy the Causal Condition, yet evolutionary biologists would shy away from considering them to be cases of natural selection. To put it otherwise, evolutionary biologists seem to isolate the causal claims that are selectively relevant from those that are not. Here are two such potential restrictions.

---

[3] Even proponents of the recent "statisticalist" view of natural selection seem to agree with this, since they avowedly do not contest statements like "variation in camouflage causes evolutionary change" (Matthen and Ariew 2009, 203).
[4] I borrow the term "Causal Condition" from Glymour (2014), but I attach a different meaning to it.

Biologists and philosophers sometimes imply (e.g. Damuth 1985; Richardson 1996) that individuals living on different continents – or separated by large geographic distances – cannot be said to belong to a population undergoing selection. Philosophers have recently begun trying to work out a set of interconnectedness criteria that a population undergoing selection should fulfill (e.g. Millstein 2010; Matthewson 2015), but the elaboration of such criteria also forces us to tackle the complicated issue of determining whether they should admit of degree (Godfrey-Smith 2009; Matthewson 2015). This restriction of the Causal Condition can be formulated by saying that a difference in phenotypic trait is said to be a *selective* (or *selectively relevant*) cause of a difference in fitness only when the set of individuals we refer to exhibit a degree of populational interconnectedness that is above a certain threshold; if that is not the case, then, even though the Causal Condition is arguably satisfied, we would still not be allowed to say that natural selection is at work in the given scenario.

But this paper is concerned with a different potential restriction of the Causal Condition, namely the idea that selective claims need to be relativized to a homogenous environment. By this, I understand that a difference in phenotypic trait is said to be a *selective* (or *selectively relevant*) cause of a difference in fitness only when the set of individuals we refer to experience a homogenous environment. This idea has been explicitly endorsed by some biologists and philosophers (e.g., again, Damuth 1985; Antonovics, Ellstrand and Brandon 1988; Brandon 1990) and probably implicitly assumed by many others. For simplicity, I will call this idea the "environmental homogeneity restriction" of the Causal Condition (hereafter, EHR). But note that other theorists do not seem to accept EHR (e.g. Glymour 2014; Abrams 2014); these authors do not claim that selective claims need not be relativized to environments; they only contest the idea that these environments must be homogenous. So, according to such authors, these environments *can* be homogenous, but they *need not* be so – it is in this sense that they reject EHR. I will not argue here for or against adopting EHR; rather, I will highlight some consequences – for single- and multi-level selection – that derive from acknowledging the distinction between adopting and rejecting EHR.

But talking about diverging positions with respect to whether selective claims should be relativized to homogenous environments requires that we first convene on what we mean by "homogeneity". Given that Brandon (1990) arguably provides the most detailed account of it, it is probably best to use his notion of environmental homogeneity here.

2.3. Brandon's notion of environmental homogeneity
Brandon (1990) distinguishes between three types of environment:

i) the external environment: the totality of abiotic and biotic factors external to the biologic entities in question;

ii) the ecological environment: the subset of the features of the external environment (as well as other factors pertaining to population structure) that affect the expected fitness of the target biological entities. The ecological environment is homogenous when the degree of adaptedness of individuals of a given type is relatively constant across the area or population. Ecological environmental homogeneity is thus defined relative to a given type: the same area may be ecologically homogenous for one type (the individuals of that type have relatively constant expected fitnesses throughout the area), and yet ecologically heterogeneous for another type (when individuals of this latter type vary in expected fitness within different sub-sections of the area, because of variation in a particular environmental feature).

iii) the selective environment: the subset of the features of the external environment (as well as other factors pertaining to population structure) to which the types represented are differentially adapted (i.e. the factors that affect the relative expected fitnesses of the types in the area or population). The selective environment is homogenous when, for each type represented, the relative adaptedness of its members is constant within the area or population (i.e., for each type, the relative expected fitness values of its members are constant). Selective environmental homogeneity is thus defined relative to all the types represented: two sub-sections of an area represent a single homogenous selective environment only if, within them, for each type represented, the relative expected fitnesses of its members are constant; if this is not the case, each sub-section becomes a selective environment in its own right.

Note that the selective environment is defined in terms of *relative* fitness, while the ecological environment is not. This is due to the fact that the ecological environment is defined relative to one type, whereas the selective environment is defined relative to all represented types. This also explains why two ecological environments can form a single selective environment and vice-versa. If only half of a field is low on nitrogen, but a resilient plant of type A can get by equally well on both halves of the field, the field as a whole will count as a single homogenous *ecological* environment relative to type A; but if, on the same field, conspecific plants of type B fare worse in the low-nitrogen half than in the nitrogen-rich half, then the two halves will count as two homogenous *selective* environments for the two types, because the *relative* fitnesses of type-B plants vary between the two halves.

Alternatively, if both types A and B fare worse in the low-nitrogen half, then the two halves will count as two (homogenous) ecological environments with respect to either type; however, if the

relative expected fitnesses of type-B (and type-A) plants do *not* vary between the two halves, then the whole field will count as a single homogenous selective environment. Otherwise put: two ecological environments are two selective environments only if selection strength is different within them. And it makes perfect sense to argue that two adjacent portions of a field constitute a single selective environment if selection strength is the same within them.

But the point to take away from this discussion is that when Brandon states that selective claims need to be relativized to homogenous environments or when other authors reject this idea, it is undoubtedly to the selective environment that they refer, since this is the only one that is defined relative to multiple types.

3. Adopting and rejecting EHR in a single-level selection scenario

In this section, I will argue that adopting or rejecting EHR makes a difference with respect to the way in which we should treat certain single-level selection scenarios. For this, I borrow a well-known example from Brandon (1990).

Imagine we have a population composed of two types of conspecific insects (G_1 and G_2) with equal initial frequencies. These insects lay their eggs on two species of plants, on which their larvae feed after hatching. The two plant species occupy two adjacent fields, which I will call environments E_1 and E_2. Now, G_1 is fitter than G_2 in any of the two environments (more of their larvae reach adulthood) because, unlike the G_2, the G_1 larvae can synthesize an enzyme that helps them metabolize a particular substance present in both plant species found in E_1 and E_2. But suppose, also, that both types are fitter in E_2 because the plant species found there constitutes better nourishment for all larvae. Here are the hypothetical fitness values assigned by Brandon:

$$W(G_1, E_1) = 1 \qquad W(G_1, E_2) = 2$$
$$W(G_2, E_1) = 0.8 \qquad W(G_2, E_2) = 1.8$$

(where $W(G_i, E_j)$ denotes the fitness of the i^{th} genotype in environment j).

But suppose also that, for reasons that remain to be specified, the two types are not distributed equally across the two environments: eggs of the fitter type G_1 are more often laid in the poor environment E_1, while G_2 eggs are more often laid in the better environment E_2. Therefore, denoting by N the number of laid eggs of each type in the first generation, and assuming that $0.8N$ of the G_1's hatch in E_1 and $0.8N$ of the G_2's hatch in E_2,[5] we obtain the following overall fitnesses:

$$W_O(G_1) = (0.8N·1 + 0.2N·2)/N = 1.2$$
$$W_O(G_2) = (0.2N·0.8 + 0.8N·1.8)/N = 1.6$$

(where $W_O(G_i)$ is the overall fitness of type i).

---

[5] To avoid unnecessary complications, I assume that all eggs hatch and that, on the whole, an equal number of eggs are laid in environments E_1 and E_2.

Even though G_1 is fitter within each environment, after one reproductive cycle it proves to be overall less successful than G_2 (less G_1 larvae reach adulthood). Brandon argues that our interpretation of this result should depend on the cause of the distribution pattern of the eggs of each type across environments.

3.1. Chance as the source of the distribution pattern

When the distribution of eggs of each type across environments is due to chance, the overall advantage of G_2 is, Brandon concludes, the result of drift, not selection. Indeed, overall, G_2 is fitter than G_1, but this is not due to their difference in traits, but to the difference between the environments E_1 and E_2 (plus the accidental unfair distribution of the two types across them). I will consequently call this Brandon's Drift-Case. However, we still have selection in this case, because the superior fitness of G_1 within any of the two environments is caused by a difference in a given trait between G_1 and G_2: the Causal Condition is thus satisfied with respect to habitat tolerance.

I agree. But I will push the analysis of Brandon's example further and show that our treatment of this case depends on the environment to which we relativize our selective claims. Thus, depending on whether or not we accept EHR, two varying analyses of the case are possible.

If, like Brandon, one relativizes selective claims to homogenous environments, then one has to:

i) note that the relative expected fitness of G_2 varies between the two environments (being 0.8 in E_1 and 0.9 in E_2). Selection strength is therefore different in E_1 and E_2 and we must conclude, with Brandon, that we are dealing with two selective environments;

ii) conclude that we have "selection occurring in two environments" here (Brandon 1990, 62), i.e. two processes of selection for habitat tolerance, one in each of E_1 and E_2, both favoring G_1.

On the contrary, if one rejects EHR, one simply concludes that selection for habitat tolerance occurs within the global environment (E_1 and E_2 put together). The difference between the two positions does not seem significant: they agree that selection for habitat tolerance occurs (and favors G_1) and that the accidental uneven distribution of types into environments favors G_2 and counterbalances the effects of selection. But the difference gets noticeable if we are interested in calculating how much of the change in the frequencies of types is due to selection for habitat tolerance.

a) If, like Brandon, one adopts EHR and, consequently, the idea that we are dealing here with "selection occurring in two environments", then calculating the change in the frequency of types that is due to selection for habitat tolerance must be done by assuming that the difference in the focal trait between the two types is to be considered a *selective* cause of their difference in fitness only when we refer to the individuals that are found within each of the two homogenous

environments. Therefore, the effects of selection for habitat tolerance need to be calculated within each environment. One way to do this is by applying Price's (1970) equation for single-level selection within each environment. (Of course, the two techniques being mathematically equivalent, one can also apply simple regression analysis to each environment.) Having determined the effects of each selection process within each environment, we need to move on and determine what this entails with respect to the overall change in the frequency of types. To do this, we take the results yielded by Price's equation within each environment and weigh them by initial within-environment population size and we further divide them by the value of the average individual fitness within the global population. Otherwise put, denoting by z the trait with respect to which genotypes differ (and supposing that like begets like in our population), the change in the population mean of z that is due to selection for habitat tolerance ($\Delta \overline{z_{ST}}$) is:

$$\Delta \overline{z_{ST}} = \frac{E(Cov_{Ej}(z_i, w_i))}{\overline{w}} \qquad \text{(Equation 1)}$$

where $Cov_{Ej}(z_i, w_i)$ is the covariance, within the j^{th} environment, between the character and fitness of the i^{th} individual, and $\overline{w}$ is the average individual fitness within the global population.

In our example, because we assumed that an equal number of eggs are, on the whole, laid in E_1 and E_2, Equation (1) becomes:

$$\Delta \overline{z_{ST}} = \frac{1}{2} \cdot \frac{\overbrace{Cov_{E1}(z_i, w_i)}^{\substack{selection \\ in \\ E1}} + \overbrace{Cov_{E2}(z_i, w_i)}^{\substack{selection \\ in \\ E2}}}{\overline{w}}$$

b) On the other hand, if one rejects EHR, one can simply assume that the difference in z is a selectively relevant cause of the difference in w within the global environment (without bothering therefore to relativize selective claims to homogenous environments). Under such an assumption, one of the analytic tools we can use for calculating the change in z that is due to selection for habitat tolerance is multiple regression analysis (see Pedhazur 1997). The idea here is to treat the fitness of each individual as depending on two independent variables, namely its phenotypic character z (i.e. whether the individual is G_1 or G_2) and the relevant property e of the environment it finds itself in (i.e. whether the individual is found in E_2 or in the less nourishing environment E_1). The following equation is employed to this end:

$$w_i = \alpha + \beta_1 z_i + \beta_2 e_i + \varepsilon_i \qquad \text{(Equation 2)}$$

Here, α is the intercept, β_1 is the partial regression of fitness on phenotypic character, controlling for the relevant environmental property, β_2 is the partial regression of fitness on the environmental property, controlling for phenotypic character, while ε_i is the unexplained residual.

We can then go on and decompose the change in the population mean of *z* in two components – the change that is due to selection on habitat tolerance and the change that is due to the uneven distribution of types into environments – as follows:

$$\Delta\bar{z} = \overbrace{\frac{\beta_1 Var(z_i)}{\overline{w}}}^{\substack{\text{selection} \\ \text{on} \\ \text{habitat} \\ \text{tolerance}}} + \overbrace{\frac{\beta_2 Cov(z_i, e_i)}{\overline{w}}}^{\text{drift}} \qquad \text{(Equation 3)}$$

where $Var(z_i)$ is the variance in *z* within the global population and $Cov(z_i, E_j)$ measures the departure from the mean in the distribution of types into environments (i.e. the extent to which one type is found more often in the richer environment).

Following Heisler and Damuth (1987), we can say that β_1 measures the expected change in fitness of an individual (controlling for the relevant environmental property) if its phenotypic character were changed by one unit. Assuming, for simplicity, that $z = 1$ for G_1 and $z = 0$ for G_2, this translates to saying that β_1 denotes the difference in fitness between the two types due to their phenotypic character *alone* (i.e. mathematically correcting for any differences in the relevant property of the environment they find themselves in). But, crucially, note that $Var(z_i)$ is calculated over the global environment (E_1 and E_2 put together). Consequently, the component $\frac{\beta_1 Var(z_i)}{\overline{w}}$ of Equation (3) measures the change in *z* that is due to differences in phenotypic character *within the global environment* (i.e. when the differences in phenotypic character are deemed selectively relevant within the global environment, and not just within each of its homogenous sections). This point is easily missed, but it cannot be avoided: someone considering that selection for habitat tolerance occurs within each of the two environments would have no choice but to apply Price's equation or some mathematically equivalent technique like *simple* regression analysis within each environment, but *not* multiple regression analysis.

Note also that the second component on the right hand side of Equation (3) is taken to represent the effects of drift on the change in *z*. The reasoning behind this is the following: one type is preferentially associated with the high quality environment (i.e. $Cov(z_i, E_j) \neq 0$). But this association, as Brandon reasoned, only reflects an accidental event, not a difference in traits between the two types. Consequently, this component is taken to reflect the effects of drift. Of course, claiming to measure the effects of drift with such accuracy sounds peculiar, but we only make this claim because we have stipulated the exact causes of fitness differences in this example. In any empirical case such a claim would be naïve, at best.

A couple of remarks need to be made at this point. It might seem odd that I have so closely linked both Price's equation and multiple regression analysis to causal claims. After all, these are

just statistical tools and inferring causation from correlation is never a good idea. However, I am not inferring causation from correlation here. In Brandon's case, the causal relations are stipulated, and the two treatments of the case yield different values for the effects of direct selection on z only because they disagree about the environments within which these causal relations are to be considered selectively relevant. Of course, when we move on to empirical cases, we would have to take additional steps (e.g. measurements on multiple correlated or uncorrelated traits, extending the measured sample, setting up control protocols and replicate populations etc.) before we could make at least mildly plausible claims about the traits or factors causing the differential fitness. But once these causal factors are plausibly identified, we can plug the variables corresponding to them into the above equations and thus obtain an estimation of the contribution of selection on the relevant trait. The point is that, just like any other equations, Equations (1)-(3) are only as causally relevant as the variables we plug into them. In the case discussed here, if differences in z and e do indeed cause differences in w, then Equations (1)-(3) are causally relevant, and choosing either Equation (1) or Equation (3) in order to estimate the change in z that is due to direct selection on z depends on whether one adopts or rejects EHR. But the reverse also holds: if these factors turned out not to be responsible for the differential fitness, then both the Pricean and the multiple regression treatments of the case would lose their causal relevance.

Let me add that, as will become obvious later (see footnote 21 below), multiple regression and Price's equation yield contrasting measurements of the change in z due to selection for habitat tolerance for all the cases in which types that differ in fitness due to their phenotypic character are, by chance, unevenly distributed across environments of different quality.

3.2. Habitat preferences as the source of the distribution pattern

We now suppose that, instead of being accidental, the distribution of types into environments was due to the behavioral preferences of the G_2 type (G_2 females tending to lay eggs in the more hospitable environment E_2 [6]). If the uneven distribution of types into environments is due to habitat preferences, then, Brandon argues, "we have a single homogenous selective environment. In this selective environment the genotype showing a preference for the richer environment has the higher relative (expected) fitness" (Brandon 1990, 63).

Brandon's conclusion is compelling. Overall G_2 is again fitter, but this time the difference in fitness between types is caused by the difference in their habitat preferences. The Causal Condition is satisfied, and we are therefore dealing with selection on habitat preferences here (so, I will call this Brandon's Selection-Case). But we also need to conclude that we are dealing with a *single* homogenous selective environment here. To see this, the question we have to ask is: to what

---

[6] Because they are conditioned to do so by some factor, be it genetic or not.

environment or to what feature of the environment is G_2 better adapted than G_1? Since G_2 is at an advantage because it is better at choosing the environment in which to lay eggs, the necessary conclusion is that it is not adapted to either E_1 or E_2, but to the *patchiness* of the global environment, i.e. to the fact that the context offers, as alternative sites for egg-laying, both environments E_1 and E_2. Therefore, the entire field, comprising both E_1 and E_2, is the relevant (homogenous) selective environment. This means that here adopting or rejecting EHR amounts to the same thing, since both positions would entail that the causes of differential fitness are to be considered selectively relevant within the global environment.[7] Therefore, irrespective of whether one adopts or rejects EHR, isolating the change in z that is due to selection on habitat tolerance from the change due to selection on habitat preferences should yield the same results. But how should this calculation be done?

It needs to be done by using multiple regression analysis. This is due to a significant difference between Price's equation and multiple regression analysis that I have not yet brought into the discussion. As Frank (2012, 1013) argued, when selection operates on two correlated traits within a single environment, Price's equation no longer allows us to partition the contributions of these processes to the overall change in z – and Frank rightly notes that in such cases Price's equation can no longer be claimed to be causally relevant and should thus be seen as nothing more than an "abstract placeholder" that can be refined by "partitioning causes by statistical regression models" (like multiple regression analysis).[8] Brandon's Selection-Case is precisely a case of this sort because, within the single homogenous environment of this case, selection acts on two correlated traits (habitat preference and habitat tolerance). They are correlated because, for a given individual, its belonging to the type with lower habitat tolerance increases its chances of being found in the richer environment; alternatively, belonging to the type that tolerates the environment better decreases one's chances of being found in the richer environment. So here, whether one adopts or rejects EHR, one can appeal to multiple regression analysis and partition the total change

---

[7] Let me add that I am *not* claiming that a researcher that is only interested in explaining why a type tolerates one environment (say E_1) better would be wrong in studying only E_1 and in considering it as a homogenous selective environment in its own right. What I am claiming – following Brandon – is that, *if one's goal is that of providing an explanation for the total evolutionary change in this population*, one would be wrong to consider E_1 and E_2 as distinct homogenous selective environments.

[8] This entails that, in cases in which the within-environment difference in fitness has a "single-cause explanation" (Frank 2012, 1014), nothing stands in the way of considering Price's equation to be causally relevant, as Frank implicitly agrees – an example of this sort was Brandon's Drift-Case analyzed above. Moreover, this is the type of case that lies at the heart of the debate about the interpretation of MLS1 that I will deal with in the second part of this paper. Indeed, this debate revolves exclusively around MLS1 cases in which it is assumed that the within-group fitness differences between individuals are due to their difference in a single relevant trait. I will uphold this assumption here. But I will also add that, if we were to shift our attention towards cases in which within-group fitness differences do not have a "single-cause explanation", as Frank puts it, speaking of "Pricean" and "contextualist" interpretations of MLS1 would become strained, since both would actually make use of some form of multiple regression analysis. However, the distinction I make here between a position that adopts EHR and one that rejects it would remain intact in such cases as well, and this corroborates my point that the Pricean and contextualist interpretations of MLS1 need to be seen as diverging positions with respect to EHR.

in z into two components: the change due to the *direct* effect of selection on the focal trait z (i.e. selection on habitat tolerance) and the change due to the *indirect* effect of selection on a different, correlated trait (namely habitat preference). This widely used method for partitioning the effects of selection on correlated traits has been developed by Lande & Arnold (1983). In our case, the partition will read:

$$\Delta \bar{z} = \overbrace{\frac{\beta_1 Var(z_i)}{\overline{w}}}^{\substack{\text{selection} \\ \text{on} \\ \text{habitat} \\ \text{tolerance}}} + \overbrace{\frac{\beta_2 Cov(z_i, e_i)}{\overline{w}}}^{\substack{\text{selection} \\ \text{on} \\ \text{habitat} \\ \text{preferences}}} \qquad \text{(Equation 3)}$$

Importantly, note that when using multiple regression analysis we treat Brandon's Selection-Case in the exact same manner as his Drift-Case. And this is only natural since, as argued above, when we use multiple regression analysis we implicitly assume that selective claims should be relativized to the global environment, whether this environment is heterogenous (as in the Drift-Case) or homogenous (as in the Selection-Case). But there is a noticeable difference between the interpretation of Equation (3) in the Selection-Case as opposed to the one from the Drift-Case: the second term on the right hand side of Equation (3) is now taken to represent the change in z due to *selection* on habitat preferences (and not *drift*, as in the interpretation that Equation (3) received in the Drift-Case). This modification is *only* due to what is reflected by the departure from the mean in the distribution of types into environments, i.e. $Cov(z_i, e_i)$: whereas in the Drift-Case the non-zero value of this covariance reflected an accidental event, in the Selection-Case it reflects a difference in a trait of the types involved (i.e. in habitat preferences).

To sum up this section, we can say that, when chance is responsible for the uneven distribution of types into environments of different quality, EHR *does* make a difference with respect to the manner in which we calculate the change in the focal trait that is due to selection. Adopting EHR entails applying Price's equation (or some mathematically equivalent technique) within the given environments, whereas multiple regression analysis can be used when one rejects EHR. By contrast, EHR does not make a difference when the uneven distribution of types into environments is due to differences in a trait between the focal individuals, because in such cases the environments of different quality constitute, together, a single homogenous environment.

4. Moving on to multi-level scenarios: an MLS1 example

So far, I have argued that there are two general ways in which the relation between natural selection and environmental homogeneity can be conceived of and that adopting or rejecting EHR also determine two different manners in which the effects of selection on phenotypic traits should be measured in cases in which types are, by chance, unevenly distributed across environments of

13

different quality. But spelling out this distinction might also have significant consequences in other – and seeming unrelated – theoretical fields. In what follows, I will thus show that rendering explicit the distinction between adopting and rejecting EHR can help clarify the terms of a recent debate in multi-level selection theory, namely the debate regarding the way in which MLS1 cases should be interpreted.

A hypothetical case will make my discussion easier to follow. This does not affect the generality of my conclusions, given that my case is similar, in most relevant respects, to most MLS1 cases discussed in the debate analyzed here (and especially with Wilson's 1975 well-known trait-group model for the evolution of altruism).

Here is the hypothetical case.[9] Imagine a population of 200 conspecific animals of two types, large (L) and small (S). The large individuals are better than the small ones at defending themselves. Large size is therefore a *directly* beneficial trait for the individual that bears it: we thus note with D_L the fitness advantage that an L-individual gets over an S-individual simply because of its own size. But suppose that, in defending itself, a large individual – but not a small one – also indirectly helps the individuals in its vicinity (because, e.g., its increased ability to resist the predator gives the other members of its group time to find shelter). Therefore, large size is also *indirectly* beneficial for the bearer's neighbors (but not for the bearer itself, who can only profit *directly* from its own defending abilities[10]). Large size therefore has direct effects on the fitness of its bearer and indirect effects on the fitnesses of the bearer's neighbors. We note with I_L the total indirect effect on fitnesses brought about by the presence of one L-individual and we also assume that the "neighbors" of an individual are all the other members of its group. This means that the indirect benefit I_L brought about by an L-individual is equally divided among all (and only) the other members of its group.

We note with B the base fitness that all individuals have to begin with (irrespective of their own or their neighbors' size), with N_j the size of the j^{th} group and with x_j the proportion of L-individuals within the j^{th} group. The fitness functions of the two types will be:

$$w_S = B + \frac{I_L x_j N_j}{N_j - 1}$$

$$w_L = B + D_L + \frac{I_L (x_j N_j - 1)}{N_j - 1}$$

The denominator ($N_j - 1$) indicates that the indirect effect of an individual's large size is shared by the *other* individuals of the group (i.e. by all N individuals of the group minus the focal

[9] This example is similar to the one I used in Jeler (2015), but the conclusions I will draw from it differ from and even contradict the ones from that paper (see footnote 16 below).
[10] I will loosen this assumption in Section 6.2.

large individual). Note that the indirect benefit received by the two types within the same group is not the same: within a given group G, an S-individual has $x_G N_G$ large neighbors, whereas an L individual has only $(x_G N_G - 1)$ large neighbors. And it is the size of one's neighbors that matters when we want to determine the indirect effects of size on fitness.

Now imagine that our metapopulation is divided into two groups O and P of size 100 each, with group O containing 20 Ls and 80 Ss, and group P containing 80 Ls and 20 Ss. For my analysis, it is unimportant whether we assume that the groups of the metapopulation remain separated for a number of generations and then fuse, before the individuals get separated again into new groups or whether the metapopulation is divided into groups only for a stage of the life cycle of the individuals (as in Wilson 1975). However, if we were to choose the latter scenario, we would also need to follow Wilson (and others, e.g. Okasha 2005) in assuming that it is only during that particular stage of their life cycle that our individuals are preyed upon by the predator that my example invokes. This indicates why ephemeral entities like groups O and P could potentially be considered units of selection: the groups might only last for part of the life cycle of individuals, but the ecological factors to which these groups could be said to be differentially adapted to also only act during the period when individuals are clustered into groups. Thus, in theory at least, groups – just like larvae, for example – can be considered differentially adapted to environmental conditions (e.g. predation), even though they are ephemeral.

Finally, assuming that $B = 9$, $D_L = 1$, $I_L = 10$, the fitnesses of types within each group become:[11]

$$w_{SO} = 11.02 \qquad w_{SP} = 17.08$$
$$w_{LO} = 11.92 \qquad w_{LP} = 17.98$$

If, for simplicity, we also assume that individuals breed truly, after one reproductive round the frequency of the L-type within the metapopulation increases from 0.5 to approximately 0.58. The debate I will discuss here, sparked by a series of papers and a book by Samir Okasha (2004a, b, 2006), concerns the way in which we should understand cases of this type. I now turn to the two competing interpretations.

5. The two interpretations as diverging positions with respect to EHR

The Pricean and the contextualist interpretations are usually seen as two competing ways of understanding MLS1 group selection (Okasha 2004a, b, 2005, 2006, 2011; Kerr 2009; Sober 2011, 2015; Goodnight 2015; Earnshaw 2015; McLoone 2015). In what follows, I will argue that it is

---

[11] The fitness values are rounded to the second decimal point, as will also be the case below for the values of frequency changes and relative fitnesses.

more accurate and more enlightening to regard them as two diverging positions with respect to EHR in MLS1 scenarios.

To understand why, let us take a closer look at my MLS1 example. I have stipulated that differences in size of individuals cause differences in fitness between them: the Causal Condition is thus satisfied and selection for large individual size is certainly at work here. But note that the relative fitness of the small type in group O is approximately 0.92, whereas in group P it is approximately 0.95: Brandon's criterion for determining whether we are dealing with multiple homogenous environments is thus satisfied by groups O and P. Therefore, my MLS1 example seems to differ with respect to Brandon's Case above only in that it is not a feature of the external environment that determines the difference in the quality of environments (like the nourishing properties of plants from environments E_1 and E_2), but a factor pertaining to population structure, namely group composition. However, recall that Brandon's criterion of relative fitness may be overridden in cases in which types are unevenly distributed into environments of different quality not by accident, but due to a trait of the types. But my MLS1 example does not provide information about the source of the distribution pattern of types into groups.

These points raise the question: to what environment or environments should we relativize the claims about selection for large individual size in a case of this sort? Three possibilities are directly relevant for us here.

a) As my analysis of Brandon's Case has already shown, someone who rejects EHR would immediately claim that these claims about selection for large individual size should be relativized to the global environment comprising both groups O and P, and this irrespective of whether or not these two groups constitute different homogenous environments. This is, I will argue below, precisely the position adopted by the contextualist interpretation.

b) On the other hand, someone who simply adopts EHR – without making any other assumptions regarding the specificity of MLS1 cases – will argue that claims about selection for large individual size should be relativized to each of the two groups *if* the difference in group composition (i.e. the unevenness in the distribution of types into groups) is due to chance; however, these claims should be relativized to the global environment comprising both groups *if* the uneven distribution of types into groups is due to a trait of types.

c) Thirdly, someone could adopt EHR while also making some other assumptions about the specificity of MLS1. This is, I will argue, the case for the Pricean interpretation. It consists in arguing that claims about selection for the focal trait should be relativized to homogenous environments, but in *also* claiming that groups varying in composition always constitute different homogenous environments *as long as we admit that, in MLS1 scenarios, groups (and not just individuals) can be seen as units undergoing selection.*

16

5.1. The Pricean interpretation

This way of interpreting MLS1 scenarios is based on Price (1972) and it has been embraced, among others, by D.S. Wilson (1975, 1979), Hamilton (1975), Wade (1985), Sober (2011, 2015), Sober and Wilson (1998). These authors usually claim that, in MLS1 scenarios, there are two kinds of entities or units undergoing selection (i.e. individuals and groups),[12] and that individual selection has to be seen as within-group selection, whereas group selection has to be seen as among-group selection. I do not think this way of presenting the Pricean interpretation is wrong, but I believe it needs to be further clarified if this interpretation is to withstand the criticism that has been traditionally thrown against it.

Using my MLS1 example, let us begin by focusing on the first claim, namely that both individuals and groups are entities or units undergoing selection in MLS1 scenarios. That there is selection for large individual size here is unquestionable, as already noted above. What about groups? Groups do differ in their composition, i.e. in the proportion of L-individuals they contain. But, in order to see whether the Causal Condition is satisfied at the group level, we need to determine whether the difference in this group-level trait causes a difference in group fitness. This obviously depends on how we define a group's fitness. So how do Priceans define a group's fitness?

To find out, we look at Price's hierarchical equation, on which this interpretation is based. According to Price (1972), we can estimate the contributions of individual and, respectively, group selection to the change in frequency of the focal trait z in the metapopulation as follows:

$$\Delta \bar{z} = \overbrace{\frac{E_j\left[Cov(z_{ij}, w_{ij})\right]}{\bar{w}}}^{\substack{\textit{individual} \\ \textit{selection}}} + \overbrace{\frac{Cov(Z_j, W_j)}{\bar{w}}}^{\substack{\textit{group} \\ \textit{selection}}} \qquad \text{(Equation 5)}$$

(where z_{ij} and w_{ij} are the character value and, respectively, the fitness of the i^{th} individual in the j^{th} group, Z_j is the group character value of the j^{th} group – i.e. the average character of its individual members –, and W_j is the mean individual fitness in the j^{th} group).

The second term on the right hand side of Equation (5) shows that, in the Pricean interpretation, a group's fitness is seen as the average individual fitness of its members. $Cov(Z_j, W_j)$ can be regarded as measuring the effects of group selection only by assuming that W_j denotes the fitness of the j^{th} group. Now, going back to my MLS1 case, does the difference in the proportion of L-individuals they contain cause the difference in fitness between groups? Yes. As

---

[12] By this I do not mean to say that, according to these authors, the Pricean interpretation necessarily sees groups as strongly identified or highly cohesive entities.

long as one accepts that it is legitimate to define a group's fitness as the average fitness of its individual members, there is little that could deny the conclusion that the Causal Condition is satisfied at the group level. Indeed, recall Glymour's (2014) definition of a direct causal relation quoted above: intervening on the difference between groups in the proportion of L-individuals they contain does alter the difference in group fitness. So, at least under Glymour's account of causation – but certainly under other accounts as well, be they manipulationist, counterfactual etc. – group selection seen as among-group selection seems to satisfy the Causal Condition.

To sum up our discussion so far, one of the crucial assumptions of the Pricean interpretation is that a group's fitness is legitimately defined as the average fitness of its individual members. Once this assumption is granted, not much seems to stand in the way of the Pricean idea that groups, and not just individuals, may be seen as units undergoing selection in MLS1 scenarios.

We now move on to the other element of the way in which the Pricean interpretation is usually presented, namely the claim that, in MLS1 scenarios, group selection should be seen as among-group selection, whereas individual selection should be seen as within-group selection. Let us begin with the latter point. That individual selection should, according to the Pricean interpretation, be seen as within-group selection is underscored by the first term on the right hand side of Equation (5), which, in fact, brings nothing new with respect to Equation (1) above, i.e. with respect to the way in which we calculated the change in *z* due to habitat tolerance in Brandon's Drift-Case when EHR was adopted. Here too, in order to calculate the contribution of individual selection to the change in *z*, we calculate the covariance between individual character and fitness within each group, and then we weigh the results by group size and further divide the result by the average fitness within the metapopulation. Therefore, the Pricean interpretation relativizes claims about individual selection to each of the groups of the focal case. From this, it obviously follows that someone who uses Equation (5) when analyzing an MLS1 case implicitly assumes that claims about selection for the relevant individual trait *need not* be relativized to the global environment of the metapopulation, i.e. to the environment containing all the groups of the case. But recall that someone who rejects EHR would immediately relativize claims about selection for the relevant individual trait to the global environment of the metapopulation. Therefore, we can say – in a provisional and somewhat complicated manner – that the Pricean interpretation rejects the rejection of EHR.

But does this mean that the Pricean interpretation adopts EHR? Does this mean that a supporter of the Pricean interpretation always relativizes claims about individual selection to homogenous environments? Given that a Pricean always relativizes claims about individual selection to each group varying in composition of the case at hand, the question becomes: do groups varying in composition always constitute different homogenous environments in the type of cases at

issue here? Yes, they do, *as long as one accepts the other main element of the Pricean interpretation that we have already discussed, namely the idea that groups may be seen as units or entities undergoing selection.*

Here is why. Recall my above discussion of Brandon's notion of environmental homogeneity: if the relative fitness of a given type varies between two sub-sections of an environment, each of these sub-sections will be considered a homogenous environment. However, the criterion of varying relative fitness can be overridden in cases in which types are unevenly distributed into these sub-sections and, moreover, this uneven distribution is due to a trait of types (i.e. is not accidental): in such cases, the two sub-sections put together constitute a single homogenous environment. Now, going back to my MLS1 example, our two groups certainly satisfy the varying relative fitness criterion, since the relative fitness of the small type is different in the two groups O and P. But can we conclude from here that the two groups *really are* different homogenous environments? If, for example, the types are unevenly distributed not by accident, but because the L-individuals prefer to assort with their own kind, should we not say that the two groups put together constitute a single homogenous environment? So how can one claim that these groups are different homogenous environments *irrespective of the source of the distribution pattern of types into groups*?

The answer is provided by the Pricean idea that, in MLS1 scenarios, groups varying in a given trait are to be seen as units on which selection acts. Remember that the type of case at issue here involves groups whose focal trait is just the mean of the traits of their individual members. Therefore, in such cases, the fact that groups contain different proportions of individual types may be seen from an individual-level perspective (and we thus speak of individual types being unevenly distributed into groups), or from a group-level perspective (and we thus speak of groups which vary with respect to a given trait, namely group composition). But, by embracing the idea that groups are units undergoing selection in MLS1 scenarios, the Pricean interpretation adopts a group-level perspective on the fact that groups contain different proportions of individual types – the uneven distribution of types into environments thus simply becomes the trait with respect to which groups vary (i.e. group composition). In other words, whether or not the uneven distribution of types into groups is due to a trait of *individual* types is now irrelevant, because this uneven distribution has simply become the trait with respect to which *groups* vary. And when the uneven distribution of types is seen from a group-level perspective, whether it is accidental or not becomes irrelevant. Indeed, whether an entity – at any level – has a trait by accident or not is irrelevant for declaring whether that trait is under selection: if one were to permanently paint a set of animals green and, as a result, they fare better than their conspecifics, we would still say that there is selection for

greenness at work,[13] irrespective of the non-accidental source of this color. Similarly, if a group is fitter than another one because it contains a higher proportion of a given type, whether this difference in group composition is accidental or is due to assortative preferences of individual types becomes irrelevant.

So, once one accepts the idea that groups are units undergoing selection in MLS1 scenarios, the uneven distribution of types into environments simply becomes the trait with respect to which groups vary. Consequently, as far as individuals are concerned, the only criterion we are left with for determining the homogeneity of their environment is that of their relative fitness: if individuals of the same type that belong to different groups vary in relative fitness, then these groups are different homogenous environments for the individuals they contain. But in all the MLS1 cases *in which there are indirect effects on fitnesses*, otherwise identical individuals will have different relative fitnesses when found in groups varying in composition: consequently, such groups will always constitute different homogenous environments. So, once it adopts the idea that groups are units undergoing selection in MLS1 scenarios, for the Pricean interpretation groups varying in character always constitute different homogenous environments. Therefore, when it states that claims about individual selection need to be relativized to each of the groups varying in character of the case at hand, the Pricean interpretation states that these claims should be relativized to homogenous environments.

Thus, by claiming that individual selection should be seen as within-group selection, the Pricean interpretation adopts EHR with respect to individual selection. For my goals in this paper, having shown this is enough. But I will briefly go on and ask whether the Pricean interpretation adopts EHR with respect to group selection as well. This is a trickier question, because the idea that group selection should be seen as among-group selection is neutral with respect to EHR: to simply state that selection at level X is selection between the biological entities at that level does not say anything about whether or not the selective claims regarding these entities should be relativized to homogenous environments. However, if one wants to hold on to the idea that the term $Cov(Z_j, W_j)$ of Equation (5) measures the effects of group selection, one should adopt EHR at the group level as well. Indeed if, for example, two groups varying in composition are found in *external* environments of different quality (e.g. one group occupies an area that happens to contain fewer predators), the covariance between group character and fitness will not be reasonably taken to measure the effects of group selection; on the contrary, a part – or, depending on the case, the entire value – of $Cov(Z_j, W_j)$ will actually be a measure of the effects of the fact that groups happen to inhabit environments of different quality (i.e. drift). Therefore, a supporter of the Pricean interpretation

---

[13] Though, of course, there will be no response to selection, because the green color will not be passed on to the next generation.

should adopt EHR at the group level as well, if she wants to hold on to the idea that the term $Cov(Z_j, W_j)$ of Price's hierarchical equation measures the effects of group selection.

But imagine that a Pricean refuses to do so, and would consequently be willing to accept Frank's (2012) suggestion that $Cov(Z_j, W_j)$ be seen as a simple abstract placeholder, and that we would have to use multiple regression analysis in order to partition the causes of the difference in fitness between groups. Whether such a position could still be called Pricean is debatable, but let us not dwell on this issue here. However, such a position would quickly prove very difficult – if not outright impossible – to defend. Here is why. Recall that, in the scenarios at issue, the character and fitness of a group are seen as the average character and, respectively, the average fitness of its individual members. A Pricean arguing that claims about among-group selection need *not* be relativized to homogenous environments would thus claim that, in cases with groups inhabiting non-homogenous environments, we are justified to look at *all* the individuals of two groups and determine to what extent the difference in (average) fitness between them is due to their (average) characters, and to what extent it is due to the difference in quality between their external environments. However, in the same case, we would not be allowed to look at *one* individual from each group and determine to what extent their difference in fitness is due to their character and to what extent it is due to the difference in their environments, because such a move would violate the adoption, by the Pricean interpretation, of EHR at the individual level. It is difficult to imagine how one would defend the idea that we are allowed to make one operation with respect to *all* the individuals of two groups, but not with respect to *each* individual from these groups. The only way out of this difficulty that I can think of is the adoption, by the Pricean interpretation, of EHR with respect to group selection as well. Bottom line: by claiming that individual selection is within-group selection, the Pricean interpretation actually claims that claims about individual selection should be relativized to homogenous environments, i.e. it *adopts* EHR at the individual level. However, I have also shown that the Pricean interpretation *should* also adopt EHR with respect to among-group selection for reasons of consistency.

Let us take stock. As noted at the beginning of this section, the Pricean interpretation is usually presented as resting on two pillars: the idea that, in MLS1 scenarios, we should see within-group selection as individual selection and among-group selection as group selection and the idea that two kinds of entities or units (i.e. individuals and groups) are undergoing selection here. My discussion has reformulated these two ideas and has shown that the Pricean interpretation may be seen as the conjunction of the following two assumptions:

A)   Selective claims should be relativized to homogenous environments (i.e. the adoption of EHR);

B)   In MLS1 scenarios, we can legitimately define a group's fitness as the mean fitness of its individual members.

But why is my reformulation preferable to the usual way of presenting the Pricean interpretation? Because it will allow us to show that the traditional objection against the Pricean interpretation (Nunney 1985a, b; Heisler and Damuth 1987; Goodnight et al. 1992; Okasha 2004a, b, 2005, 2006, 2011; Goodnight 2015) is unconvincing, because its two extant versions either ignore or beg the question against assumption (A). But this will be shown in section 6, after the introduction of the contextualist interpretation.

5.2. The contextualist interpretation

The contextualist interpretation appeared as a reaction to the Pricean interpretation when Heisler and Damuth (1987) proposed "contextual analysis" as a statistical tool for analyzing MLS1 cases. As they acknowledge, the variety of methods collectively known as "contextual analysis" are a form of multiple regression analysis and have been "developed within the social sciences as tools for interpreting the effects of social, institutional, and political contexts on human behavior" (583).

When this method is applied to MLS1 scenarios, the first step is to write the fitness of each individual of the focal case as a function of its own individual character z and of the character of the specific context it finds itself in. In my MLS1 example, the relevant contextual character is the proportion of L-individuals in the neighborhood of a given individual, because an individual's fitness depends not only on its own size, but also on the size of its *neighbors*. Therefore, if we define the neighborhood character h_{ij} as the average individual character of the neighbors of the i^{th} individual in the j^{th} group,[14] the fitness of each individual in my MLS1 example can be written, just like in Equation (2) above, as:

$$w_{ij} = \alpha + \beta_I z_{ij} + \beta_C h_{ij} + \varepsilon_i$$

(where β_I is the partial regression coefficient on individual character, controlling for contextual character, and β_C is the partial regression coefficient on contextual character, controlling for individual character).

Since the intercept α is the same for all the individuals of the given case and since the variance of the unexplained residual ε_i must be minimized, this decomposition allows us to express the *differences* in fitness between individuals as a function of only two factors, namely the

---

[14] See Okasha (2004b, 2005, 2006) for a discussion – based on Nunney (1985a) – of the cases in which this choice of variables is appropriate. The application of contextual analysis by using, as an independent variable, the "neighborhood character" and, respectively, the "group character" is what Okasha calls the "neighborhood" and, respectively, the "contextual approaches". But the two are essentially the same approach and only differ – very slightly – in their choice of variables (the neighborhood character being the preferable variable when an individual's fitness depends on the average character of its neighbors, whereas the group character is the preferable variable for cases in which an individual's fitness depends on the average character of its group, the focal individual included).

difference in individual character *z* and the difference between the neighborhoods individuals find themselves in. Consequently, once the partial regression coefficients are determined, we can estimate the change in *z* that is due to direct selection on *z* and the change in *z* that is due to the uneven distribution of types into neighborhoods in the same manner as in Equation (3) above:

$$\Delta \bar{z} = \overbrace{\frac{\beta_I Var(z_{ij})}{\bar{w}}}^{\substack{\text{individual} \\ \text{selection}}} + \overbrace{\frac{\beta_C Cov(z_{ij}, h_{ij})}{\bar{w}}}^{\substack{\text{evolutionary} \\ \text{change} \\ \text{due} \\ \text{to} \\ \text{uneven} \\ \text{distribution}}}$$

Note that – as highlighted by the fact that they employ the exact same equations – the contextualist interpretation sees MLS1 cases precisely like someone rejecting EHR treated Brandon's Case above. Indeed, according to the contextualist interpretation, the right way to understand MLS1 scenarios needs to start from the observation that here the average fitnesses of two (or more) types vary as a result of two factors: their difference in the focal trait *z* and the fact that, on average, types are found in contexts of different quality. In my MLS1 example, the L-individuals are at an advantage both because of their own size *and* because they tend to be found more often in the company of other L-individuals (and consequently benefit, on average, more than the S-individuals from the indirect effects of large size). Otherwise put, L-individuals are favored both by direct selection on body size and by the fact that they tend to be preferentially found in the environment containing more indirect effects of large size. Therefore, the only difference between the contextualist interpretation of MLS1 and the multiple regression treatment of Brandon's Case above (applied by someone rejecting EHR), lies in the fact that in the latter case the variable *e* referred to the environmental quality that depended on a factor of the external environment (namely the nourishing properties of the plant species occupying environments E_1 and E_2), whereas in the former case, the variable *h* refers to the environmental quality that depends on a factor pertaining to population structure (see Heisler and Damuth 1987, 583). But this difference in one of the variables employed need not obscure the main point, namely that the contextualist interpretation treats MLS1 cases precisely like someone rejecting EHR treated Brandon's Case above.

Three notable points ensue.

i) Just like when we applied multiple regression analysis in Brandon's Case – since, again, β_I and $Var(z_{ij})$ are calculated over the global environment –, here too the "individual selection" component is calculated over the global environment comprising both groups O and P. The contextualist interpretation thus assumes that, instead of relativizing claims about individual selection to each group, we should relativize them to the global environment. Otherwise put, the central assumption of the contextualist interpretation is precisely the negative assumption consisting

in the rejection of EHR. Once this assumption is made, it is straightforward to go on and isolate the effects of direct selection on *z in the global environment* from the effects of the uneven distribution of types into neighborhoods.

ii) Note that I have labeled the second component on the right hand side of the contextualist partition as "evolutionary change due to uneven distribution" (of types into neighborhoods). I have deliberately avoided labeling it as it is usually labeled in the philosophical literature, namely "selection on neighborhood character" (Okasha 2006), "selection on belonging to variables" (Glymour 2008) etc.[15] This is because, just like in Brandon's Case, whether we say that this evolutionary change is due to drift or selection depends on whether the non-zero value of $Cov(z_{ij}, h_{ij})$ – i.e. the uneven distribution of types into neighborhoods – reflects an accidental event or a phenotypic trait (e.g. assortative preferences). If this uneven distribution reflects a chance event, we need to conclude that *drift* is the second factor leading to evolutionary change here. The fact that, by chance, L-individuals are found more often in the company of other L-individuals and get a fitness advantage because of this is not a reason for us to say that selection is at work here. On the other hand, if the uneven distribution of types reflects a difference in a phenotypic trait (e.g. assortative preferences), then we need to conclude that selection for *this* trait is responsible for the fraction of the evolutionary change calculated by $\beta_C Cov(z_{ij}, h_{ij})$. Therefore, just like in Brandon's Selection-Case above, we would thus be dealing here with two processes of selection, namely selection on body size and selection on assortative preferences.[16] To sum up, the contextualist interpretation of MLS1 can only state whether the evolutionary change due to the uneven distribution of types should be considered the result of drift or of selection *if* the source of the distribution pattern of types is specified. To my knowledge, nobody has yet acknowledged this fact, and this constitutes one of the significant results that stem from the framework proposed here.

iii) But this also entails that, according to the contextualist interpretation, groups do not constitute "units of selection" in MLS1 scenarios.[17] As just noted, what is usually called "group selection" in the contextualist interpretation (i.e. selection on group/neighborhood character) can in fact only be called selection when the non-zero value of the covariance between individual character and

---

[15] Or "group selection *sensu* the contextual approach" (Okasha 2006, Earnshaw 2015) as it is called when fitnesses depend on the average frequency of the focal trait within an individual's group (and not just within an individual's neighborhood).

[16] In Jeler (2015, 2016), I argued that what the contextualist interpretation calls "group selection" (i.e. "selection on neighborhood/group character", depending on the details of the case) can be seen as a form of frequency-dependent selection in multi-group scenarios. If true, this idea would hold irrespective of whether the non-zero value of the covariance between individual character and neighborhood/group character reflects a chance event or a phenotypic trait. However, this amounts to nothing more than claiming that the *only* factor that is responsible for this "selection on neighborhood/group character" is the fact that one type is found more often in the richer environment. But, as shown above, the differential fitness that comes from the uneven distribution of types into environments of different quality cannot be called selection unless this distribution pattern reflects a difference in phenotypic trait (e.g. assortative preferences).

[17] Gardner (2015b) makes a similar point, but brings little arguments in its support.

neighborhood/group character is due to a phenotypic trait with respect to which types differ (e.g. assortative preferences). But in such situations, we are in fact dealing with selection on the assortative preferences *of individuals* (or, more generally, selection on the *individual* traits that bias the distribution of types into groups or neighborhoods). This is selection on an individual trait – hence individual selection – and not group selection. A corollary to this is that the contextualist interpretation has no use for something like the assumption (B) of the Pricean interpretation.

Most authors taking part in the debate discussed here pay no attention to these three points when they compare "the Price" and "the contextual approaches" – this is why I avoided using these terms here. However, points ii) and iii) are, in some respects, similar to an idea from Nunney (1985a), who claimed that all that can be called "group selection" in MLS1 cases is selection on the *individual* traits that lead to non-random group formation.[18] I think he is perfectly right with respect to the contextualist interpretation of MLS1, but not necessarily with respect to MLS1 in general. In order for this to be right about MLS1 in general, one would need to also prove the Pricean interpretation wrong. To this end, Nunney (1985a, b) has provided what can now be considered the "traditional" objection against the Pricean interpretation. I will end this paper by showing that this objection becomes unconvincing once the issue of EHR is brought into the picture.

6. Why the traditional objection to the Pricean interpretation is unconvincing

The traditional objection to the Pricean interpretation comes in two versions, a narrow and a wider one. I will discuss them in turn.

6.1. The narrow version

The narrow version of the objection consists in showing that Price's hierarchical equation – Equation (5) above – will yield non-zero values for the contribution of group selection to the overall change in z even in cases in which the fitnesses of individuals depend entirely on their own individual character (see Nunney 1985a; Heissler and Damuth 1987; Okasha 2004a, b, 2005). To see this more clearly, let us completely eliminate from my MLS1 example the indirect effects of large size on the fitnesses of the bearer's neighbors. The fitness functions of types within the two groups thus become:

$$w_{SO} = w_{SP} = B = 9$$

$$w_{LO} = w_{LP} = B + D_L = 10$$

The fitnesses of individuals no longer depend on the composition of their groups. However, because there are more L-individuals in group P, the mean fitness within group P becomes 9.8,

---

[18] Bourrat (2015) recently made a very similar claim, but his argument is seriously undermined by the assumption that natural selection can only act on "intrinsic-invariable properties". This is tantamount to claiming, for example, that frequency-dependent selection is not natural selection.

whereas the mean fitness in group O is 9.2, and this difference in mean fitnesses is explained by the fact that groups contain different proportions of L-individuals. In Nunney's (1985a, 219) words, an illusion of group selection thus "arises because by chance some groups contain more of the fittest genotype, and these groups are the most successful". Equation (5) will thus detect group selection here (i.e. $Cov(Z_j, W_j) \neq 0$), even though all of the fitness difference between groups is due to the fact that some groups contain a larger proportion of the type that is fitter *anyway* (i.e. the type that is fitter solely due to the direct effects of its body size). Supporters of the contextualist interpretation claim it would be wrong to see the change in *z* that is measured here by $Cov(Z_j, W_j)$ as being the result of group selection; on the contrary, they claim, it should count as the result of individual selection.

However, if my presentation of the Pricean interpretation is right, a Pricean would agree with this claim. Recall that the first assumption of the Pricean interpretation consists in the adoption of EHR. Or, in the case at hand, we are indeed dealing with a single homogenous environment: given that group membership does not affect the fitnesses of individuals, the individuals of the same type have constant relative fitnesses across the two groups. Therefore, we have to conclude that the global environment (containing both groups O and P) constitutes a single homogenous environment for all the individuals of the metapopulation. We are therefore dealing here with a single process of (individual) selection on body size, acting within the whole metapopulation. The first assumption of the Pricean interpretation thus prevents us from even applying Price's hierarchical equation to such a case. Consequently, this narrow version of the objection to the Pricean interpretation misses its target.

The narrow version has never been a knock-down argument against the Pricean interpretation. Supporters of the latter have traditionally countered this objection by imposing constraints on what counts as a group, i.e. by arguing that a group should only be considered a group if there are fitness-affecting interactions between its members or, in my terms, only if there are indirect effects of the focal trait on the bearer's neighbors (see Sober 1984, 2011; Okasha 2006; McLoone 2015; Jeler 2016). So why did I provide my own response to this objection by appealing to environmental homogeneity? Because the interactionist definition of groups does not also defuse the wider version of the objection to the Pricean interpretation, while my solution does. And the discussion of the narrow version in terms of environmental homogeneity serves as a good introduction for the discussion of the wider version.

6.2. The wider version

The narrow version of the objection consisted in pointing out that the Pricean covariance between group character and fitness will be non-zero even in cases in which *all* of the fitness

difference between groups is due to the fact that some groups contain a larger proportion of the type that is fitter anyway (i.e. even in cases *without* indirect effects of the focal trait on the fitnesses of individuals). On the other hand, the wider version consists in similarly pointing out that, even in cases in which there are indirect effects of the focal trait on the fitnesses of individuals, *some* of the Pricean covariance between group character and group fitness will again be simply due to the fact that some groups contain a larger proportion of the type that is fitter anyway.

To get a better picture of this wider objection, we go back to my MLS1 example and assume, this time, that the fitnesses of individuals depend on their own individual character and on the character of their group (not their neighborhoods). Therefore, the group (and no longer the group minus the focal individual) becomes the relevant context for an individual's fitness here. In other words, the fitness functions of the two types in the j^{th} group change slightly and become:

$$w_{Sj} = B + I_L x_j$$

$$w_{Lj} = B + D_L + I_L x_j$$

The numeric values of the fitnesses of types within groups are thus slightly altered:

$$w_{SO} = 11 \qquad w_{SP} = 17$$

$$w_{LO} = 12 \qquad w_{LP} = 18$$

Note that, again, *if we completely disregard, for a second, the indirect effects of large size on fitnesses*, the mean individual fitness between the groups still differs, being 9.2 in group O and 9.8 in P, and this is due to the fact that group P contains more L-individuals that are better at defending *themselves* (and not due to the indirect effects of their size). However, this difference in group fitness that is entirely due to the direct effects of large size *does* constitute a part of $Cov(Z_j, W_j)$, the Pricean component measuring the effects of group selection. This fraction of the covariance, the objection goes (see Nunney 1985b; Heissler and Damuth 1987; Goodnight et al. 1992; Okasha 2006; Goodnight 2015), should, again, not count as group selection, but it should count as individual selection since it is nothing more than the result of the fact that some groups contain a larger proportion of the type that is fitter anyway. Thus, Okasha (2006, 98) states that "the term *Cov(W, Z)* will *always* contain an element which is a by-product of particle-level selection"; Goodnight et al. (1992, 745) make the same objection when arguing that "the group-level covariance components can include both covariance due to the direct effects of selection at one level and indirect effects of selection at another level."

This objection is unconvincing. To see why, note – as Goodnight et al. indicate – that the covariance between group character and group fitness can be decomposed as follows:

$$Cov(Z_j, W_j) = \beta_I Var(Z_j) + \beta_C Var(Z_j) \qquad \qquad \text{(Equation 6)}^{19}$$

(where, just like above, following Heisler and Damuth 1987, the partial regression coefficients β_I and β_C measure the expected change in fitness of an individual if its individual character and, respectively, contextual character – i.e. group character, here – were changed by one unit, while everything else were held fixed).

Equation (6) clearly shows that the narrow version of the objection is just a special case of the wider version: in the narrow version, $\beta_C = 0$ (because there are no indirect effects of the focal trait on fitnesses) and therefore Equation (6) reduces to $Cov(Z_j, W_j) = \beta_I Var(Z_j)$. As seen above, the narrow version consisted in claiming that the change in z measured by the term $\beta_I Var(Z_j)$ should be seen as resulting from individual – and not group – selection. The wider version of the objection consists in saying that, in cases where there are indirect effects of the focal trait on fitnesses (i.e. when $\beta_C Var(Z_j) \neq 0$), the term $\beta_I Var(Z_j)$ should *also* be moved under the label "individual selection".

But this time, the Pricean disagrees. To see why, recall that in all MLS1 cases in which there are indirect effects of the relevant trait on fitnesses, groups varying in character will constitute different homogenous environments.[20] For example, in the case at hand, the two groups constitute two homogenous environments (the relative fitness of type-S individuals is not the same in groups O and P). Secondly, to claim that the fraction of the difference in fitness between groups that is due to the fact that some groups contain a larger proportion of the type that is fitter anyway should count as individual selection assumes that claims about direct selection for body size should be relativized to the global environment. Indeed, the litigious term $\beta_I Var(Z_j)$ measures the difference in group fitness that is due solely to the size of their individual members, *given that the groups contain different proportions of types* (recall that $Var(Z_j)$ denotes the departure from the mean in the distribution of types into groups). Claiming that $\beta_I Var(Z_j)$ should be subsumed under "individual selection" assumes that differences in size of individuals that do not belong to the same group should be considered selectively relevant for individual selection. Otherwise put, the claim made by the critiques of the Pricean interpretation assumes that differences in the focal phenotypic trait of

---

[19] Without further dwelling on this point, I will just indicate that, as Okasha (2006, 89) has shown, Equation (6) may be derived from the above discussed Equation (2) of multiple regression analysis if we keep in mind that in the MLS1 cases at issue here the fitness of a group is the average fitness of its members – and consequently $Cov(Z_j, W_j) = Cov(Z_j, w_{ij})$ – and a group's character is the average character of its individual members – and consequently $Cov(z_{ij}, Z_j) = Var(Z_j)$.

[20] Obviously, in cases in which the indirect effects of the focal trait on fitnesses are very small, the differences between the relative fitnesses of types across the given groups might be extremely small and one could reasonably discount these "grey-area" cases as involving multiple homogenous environments.

individuals should be deemed selectively relevant not within each of the groups varying in character, but within the global environment (all groups of the case put together).[21] As seen above, the Priceans agree with this *only* in cases in which group membership does not affect individual fitnesses (i.e. only in cases to which the narrow version of the objection refers); however, in all cases in which there are indirect effects of the focal trait on fitnesses, they will see groups varying in character as different homogenous environments. Therefore, to assume that in the latter type of cases claims about individual selection should also be relativized to the global environment is to beg the question against the first assumption of the Pricean interpretation, namely the adoption of EHR.[22]

Therefore, as long as the Pricean adoption of EHR goes unchallenged, the wider version of the objection is unconvincing. However, one could still argue that Equation (6) shows that the differential group fitness actually supervenes on differences in fitness between the individuals (of the two groups) that are due to *two* factors, namely direct and, respectively, indirect effects of the focal trait. But a Pricean can always reply that this is an individual-level description (Sober 2011), because the direct and indirect effects make a difference for the fitnesses of *individuals*. Indeed, it is good for an individual to be large because it is better able to resists predators; likewise, it is good for an individual to have L-individuals in its vicinity, because they delay the predator and thus allow the focal individual to flee. But, holding all other factors unchanged, what is good for a group? Simply that it contain a larger proportion of large individuals (i.e. that it have a higher group character value). The Causal Condition thus seems satisfied at the group level: a difference in a (single) group-level property seems to cause the difference in fitness of groups.

A more general approach remains open: one can always rephrase the above objection and simply note that, in MLS1 scenarios, the differential group fitness *supervenes* on differences in fitness between individuals. Given this dependence of group fitnesses on individual fitnesses, instead of looking for purported causal relations between character and fitness at the group level, the objection would go, we would be better off looking for them at the more basic individual level. But this objection is different from the traditional objection against the Pricean interpretation,

---

[21] And this is, of course, the basic assumption of the contextualist interpretation. Once we move the litigious term $\beta_I Var(Z_j)$ from "group" to "individual" selection, the Pricean partition of the overall change in z into components due to individual and group selections collapses into the contextualist partition because, as it can easily be shown, for cases in which individual and contextual characters linearly affect fitnesses,

$$\frac{E(Cov_j(z_{ij}, w_{ij}))}{\overline{w}} + \frac{\beta_I Var(Z_j)}{\overline{w}} = \frac{\beta_I Var(z_i)}{\overline{w}}.$$ This equation also holds for cases of single-level selection with accidental distribution of types across environments of different quality, like Brandon's Drift-Case from Section 3 (of course, in such cases the subscript j will refer to the different environments of the case – and not to groups –, and Z_j will thus denote the average value of the phenotypic character in the j^{th} environment).

[22] Sober (2015, 839) – followed by McLoone (2015) – suggests that the wider version of the objection "may beg the question", but he does not bring further arguments in support of this claim. In Jeler (2016), I contested Sober's point, but on grounds that become questionable once we bring EHR into the discussion.

because it consists in outright contesting the Pricean assumption (B), namely the idea that a group's fitness is legitimately defined as the average fitness of its individual members. Otherwise put, it consists in contesting the idea that groups may be seen as units of selection in MLS1 scenarios. And given that, as argued above, the contextualist interpretation has no need for assumption (B) and consists precisely in looking only at the individual level for potential selective processes at work in MLS1 cases, supporters of the contextualist interpretation should probably concentrate on contesting the Pricean assumption (B) (or, alternatively, on contesting the other main Pricean assumption, namely the adoption of EHR). I thus suggest that the debate about which of the two interpretations of MLS1 is preferable should shift its focus towards assessing the validity of the two assumptions of the Pricean interpretation.

Conclusion

This paper has identified two general positions with respect to the relationship between environment and natural selection. They consist in claiming that selective claims need and, respectively, need not be relativized to homogenous environments. Distinguishing in this manner between the adoption and the rejection of EHR may cast a new light on some disputed or insufficiently clarified points in evolutionary theory, but I have focused here on showing how it can help clarify the terms of a recent debate in multi-level selection theory. More precisely, I have shown that, of the two competing interpretations of MLS1 – namely the Pricean and the contextualist interpretations –, the former adopts EHR, while the latter rejects it. A few significant points derive from this. First of all, contrary to what is usually assumed in the literature, according to the contextualist interpretation, only individuals may be said to be under selection in MLS1 scenarios. But – and, to my knowledge, this is a completely novel claim in the field – I have also shown that the contextualist interpretation is sensitive to information about the source of the distribution of individual types into groups: according to this interpretation, depending on whether the distribution of types into groups is random or not, MLS1 scenarios either involve just individual selection and drift, or they involve only selection on two correlated individual traits (one of these being the individual trait that leads to non-random distribution of types into groups).

Therefore, of the two interpretations, only the Pricean one is a properly multi-level selectionist interpretation, i.e. one that allows groups to function as "units of selection" in MLS1 scenarios. The debate between supporters of the two interpretations should thus be seen as a debate about whether or not MLS1 scenarios are scenarios in which selection may be said to operate at multiple levels. On this front, seeing the Pricean and the contextualist interpretations as diverging positions with respect to EHR has allowed me to show that the main traditional objection against the Pricean interpretation is unconvincing. Consequently, I suggest that opponents of the Pricean

interpretation (and, implicitly, of the idea that group selection is at work in MLS1 scenarios) should shift their focus and either contest, on general grounds, EHR or argue against the other main assumption of the Pricean interpretation, namely the idea that a group's fitness is legitimately defined as the average fitness of its individual members.

Acknowledgements

References

Abrams M (2014) Environmental grain, organism fitness and type fitness. In: Barker G, Desjardin E, Pearce T (eds) Entangled life. Organism and environment in the biological and social sciences. Springer, Dordrecht, pp 127–151

Antonovics J, Ellstrand NC, Brandon R (1988) Genetic variation and environmental variation: expectations and experiments. In: Gottlieb LD, Jain SK (eds) Plant evolutionary biology. Chapman & Hall, London, pp 275–303

Bourrat P (2015) Distinguishing Natural Selection from Other Evolutionary Processes in the Evolution of Altruism. Biol Theory 10:311–321

Brandon RN (1990) Adaptation and Environment. Princeton University Press, Princeton

Damuth J (1985) Selection among 'species': a formulation in terms of natural functional units. Evolution 39:1132–1146

Damuth J, Heisler L (1988) Alternative Formulations of Multilevel Selection. Biol Philos 3:407–430

Earnshaw E (2015) Group selection and contextual analysis. Synthese 192:305–316

Frank SA (2012) Natural selection. IV. The Price equation. J Evol Biol 25:1002–1019

Gardner A (2015a) The genetical theory of multilevel selection. J Evol Biol 28:305–319

Gardner A (2015b) More on the genetical theory of multilevel selection. J Evol Biol 28:1747–1751

Glymour B (2008) Correlated Interaction and Group Selection. Br J Philos Sci 59:835–855

Glymour B (2014) Adaptation, adaptation to, and interactive causes. In: Barker G, Desjardin E, Pearce T (eds) Entangled life. Organism and environment in the biological and social sciences. Springer, Dordrecht, pp 105–126

Godfrey-Smith P (2009) Darwinian populations and natural selection. Oxford University Press, Oxford

Goodnight CJ (2015) Multilevel selection theory and evidence: a critique of Gardner, 2015. J Evol Biol 28:1734–1746

Goodnight CJ, Schwartz JM, Stevens L (1992) Contextual analysis of models of group selection, soft selection, hard selection and the evolution of altruism. Am Nat 140:743–761

Hamilton WD (1975) Innate social aptitudes of man: An approach from evolutionary genetics. In: Fox R (ed) Biosocial anthropology. Wiley, New York, pp 133–155

Heisler IL, Damuth J (1987) A method for analyzing selection in hierarchically structured populations. Am Nat 130:582–602

Jeler C (2015) Is there such a thing as "group selection" in the contextual analysis framework? Hist Philos Life Sci 36:484–502

Jeler C (2016) Do we need a new account of group selection? A reply to McLoone. Biol Theory 11:57–68

Kerr B (2009) Theoretical and experimental approaches to the evolution of altruism and the levels of selection. In: Garland Jr T, Rose MR (eds) Experimental evolution. Concepts, methods, and applications of selection experiments. University of California Press, Berkeley, pp 585–630

Lande R, Arnold SJ (1983) The measurement of selection on correlated characters. Evolution 37:1210–1226

Matthen M, Ariew A (2009) Selection and Causation. Philos Sci 76:201–224

Matthewson J (2015) Defining paradigm Darwinian populations. Philos Sci 82:178–197

McLoone B (2015) Some criticism of the contextual approach, and a few proposals. Biol Theory 10:116–124

Millstein R (2010) The concepts of population and metapopulation in evolutionary biology and ecology. In: Bell MA, Futuyama DJ, Eanes WF, Levinton S (eds) Evolution since Darwin: The first 150 years. Sinauer, Sunderland, pp 61–85

Nunney L (1985a) Group selection, altruism, and structured-deme models. Am Nat 126:212–230

Nunney L (1985b) Female-biased sex ratios: individual or group selection?. Evolution 39:349–361

Okasha S (2004a) Multi-level selection, covariance and contextual analysis. Br J Philos Sci 55:481–504

Okasha S (2004b) Multilevel selection and the partitioning of covariance: a comparison of three approaches. Evolution 58:486–494

Okasha S (2005) Altruism, group selection and correlated interactions. Br J Philos Sci 56:703–725

Okasha S (2006) Evolution and the levels of selection. Oxford University Press, Oxford

Okasha S (2011) Reply to Sober and Waters. Philos Phenomen Res 82:241–248

Pedhazur EJ (1997) Multiple regression in behavioral research. Explanation and prediction. Thomson Learning, Boston

Price G (1970) Selection and covariance. Nature 227:520–521

Price G (1972) Extension of covariance selection mathematics. Ann Hum Gen 35:485–490

Richardson R (1996) Critical notice: Robert Brandon, *Adaptation and Environment*. Philos Sci
63:122–136

Sober E (1984) The nature of selection: Evolutionary theory in philosophical focus. MIT Press,
Cambridge

Sober E (2011) Realism, conventionalism, and causal decomposition in units of selection:
reflections on Samir Okasha's *Evolution and the Levels of Selection*. Philos Phenomen Res
82:221–231

Sober E (2015) Replies to commentators on *Did Darwin Write the Origin Backwards?*. Philos Stud
172:829–840

Sober E, Wilson DS (1998) Unto others: the evolution and psychology of unselfish behavior.
Harvard University Press, Cambridge

Wade MJ (1985) Soft selection, hard selection, kin selection, and group selection. Am Nat 125:61–
73

Wilson DS (1975) A theory of group selection. PNAS 72:143–146

Wilson DS (1979) Structured demes and trait-group variation. Am Nat 113:606–610