

Gestalt-Switching and the Evolutionary Transitions

Peter Godfrey-Smith and Benjamin Kerr

ABSTRACT

Formal methods developed for modeling levels of selection problems have recently been applied to the investigation of major evolutionary transitions. We discuss two new tools of this kind. First, the ‘near-variant test’ can be used to compare the causal adequacy of predictively equivalent representations. Second, ‘state-variable gestalt-switching’ can be used to gain a useful dual perspective on evolutionary processes that involve both higher and lower level populations.

- 1 *Introduction*
 - 2 *Summary of our Modeling Framework*
 - 3 *The Near-Variant Test for Causal Adequacy*
 - 4 *State-variable Gestalt-Switching*
 - 5 *Representing an Evolutionary Transition*
 - 6 *Conclusion*
-

1 Introduction

The ‘major transitions’ in evolution are a central topic in recent evolutionary theory, and a rapidly emerging one in philosophy of biology. One consequence of this work has been a transformation of debates about the ‘levels’ or ‘units’ of selection. Earlier treatments of those problems assumed the existence of the familiar biological hierarchy, and asked where in this hierarchy selection should be seen as operating. The evolutionary transitions, however, are the processes in which new levels of the familiar hierarchy came into being (Buss [1987]; Maynard Smith and Szathmáry [1995]; Michod [1999]). The levels of selection debate has thus undergone its own transition, from a solely *synchronic* to a partly *diachronic* orientation (Okasha [2005]). The new questions become: How do evolutionary processes acting at the level of independent lower level entities produce new biological individuals visible

at a higher level? What marks the appearance of a genuine higher level individual? And—the topic of this article—how might such ‘transitions in individuality’ be best represented in formal models?

In earlier work we looked at the role of two alternative ‘perspectives’ on a particular kind of evolutionary model (Kerr and Godfrey-Smith [2002a]). This is the ‘trait group’ model (Wilson [1980]), which features a cycle of formation and dissolution of groups, systematically related to the life cycle of the lower level individuals or ‘particles’ that make up the groups. Following a tradition of pluralist work (Dugatkin and Reeve [1994]; Sterelny [1996]), we showed that a model of this kind can be *parameterized* in two ways—an individualist way, and a multi-level way. We gave a ‘translation manual’ that shows the mathematical equivalence of the two frameworks, and argued for *gestalt-switching pluralism*. The two frameworks are formally equivalent, but each perspective ‘packages’ information differently and has different heuristic features. There are positive advantages to switching back and forth between the two.

Here we look at extensions of that work, with a particular focus on evolutionary transitions. The extension involves two tools. One is a test for the *causal adequacy* of a representation. Suppose we have two descriptions of a process that are both predictively adequate. Might one nonetheless be more faithful than the other to the causal structure of the system? We argue that this question can be assessed by investigating how the two alternative descriptions must be modified to deal with *near-variants* of the original process. The second is the demonstration of the possibility of a new kind of gestalt-switching. Rather than moving between alternative sets of parameters which represent the fitnesses of entities that figure in the process, we now consider a switch between sets of *variables* that track the frequencies of different entities. We argue that this second kind of gestalt-switching has special relevance to evolutionary transitions.

2 Summary of our Modeling Framework

Our models assume a low-level population of ‘particles’, of types *A* and *B*, which are collected at some point in their life cycle into temporary groups of fixed initial size (*n*). Reproduction is asexual and faithfully preserves a particle’s type. After reproduction the groups dissolve, yielding a new pool of particles that start the cycle anew.

A model of change in such a system will feature equations that make use of three ingredients (see also Table 1):

- (i) a set of *state variables* that represent the frequencies of particle types,
- (ii) a *fitness structure* that specifies the fitness properties of individuals, groups, or both, and

Table 1. Basic components of models of selection in group-structured populations

Model component	Description
State variables	Variables tracking the frequencies of different particles
Fitness structure	Parameters giving fitnesses of particles and/or groups
Group frequency distribution	Distribution describing the proportions of different groups

- (iii) a *group frequency distribution*, which specifies, for time-step t , the frequency $f_i(t)$ of groups containing i members of the A -type (and hence $n - i$ members of the B -type).

In our earlier work ([2002a]), ‘gestalt-switching’ involves moving between two ways of formulating the fitness structure, as it relates to the role of groups. We may use:

- (i) A *contextual* treatment of group structure, yielding an *individualist* description of natural selection. Here the A -type is associated with a set of ‘ α ’ parameters, where α_i is the absolute fitness of a particle of the A -type in a group containing a total of i particles of the A -type (including itself) and $n - i$ particles of the B -type. For the B -type, we have a corresponding set of ‘ β ’ parameters, where β_i gives the fitness of a B -particle in a group with i A particles.
- (ii) A *collective* treatment of group structure, yielding a *multi-level* description of natural selection.¹ Now a set of ‘ π ’ parameters represent the combined outputs of various groups. Specifically, π_i is the total productivity from a group containing i A -types. A set of ‘ ϕ ’ parameters specifies how this productivity is distributed between particle types. Here ϕ_i is the proportion of π_i that is of the A -type.

Representations of change in such a system can be translated from one form to the other without loss of information. For example: $\pi_i = i\alpha_i + (n - i)\beta_i$; $\phi_i = i\alpha_i / (i\alpha_i + (n - i)\beta_i)$. However, each parameterization ‘packages’ information differently and involves different applications of some core Darwinian concepts. The translation is not akin to a metric-to-imperial conversion in which all theoretical concepts retain the same role. One framework encourages us to think in terms of autonomous individuals interacting within a social context; groups, in such a framework, are not bearers of fitness or competing entities. The other framework encourages us to recognize a set of higher level collectives as entities that compete and differ in fitness.

¹ In the terminology of Damuth and Heisler [1988], this is a multi-level representation in the ‘MLS1’ sense, as opposed to ‘MLS2’. The relations between MLS1 and MLS2 frameworks will be discussed below.

3 The Near-Variant Test for Causal Adequacy

The first tool we introduce is a test for how faithful a representation of a process is to the *causal structure* of that process, even where we assume the representation is predictively adequate. We approach these questions via what we call ‘near-variant analysis’. This framework is related to a family of recent approaches to causal questions that emphasize manipulation and intervention (Pearl [2000]; Woodward [2003]), and also to some non-interventionist counterfactual approaches (Lewis [2000]).

The core idea is as follows. Assume we have two formally adequate descriptions of a system undergoing change, where each description packages information differently. To the extent that a description is faithful to the system’s causal structure, the following should hold: when we consider a *near-variant* of the system under consideration, a system of the same kind but with some slight modification, we will be able to construct a representation of the new system via only slight modifications of the existing representation. Most intuitively, if we consider a near-variant that involves a *localized* change to the first system’s structure, it ought to be possible to construct an adequate representation of the new system via a localized change to the old representation. In this second sense, a ‘localized’ change is a modification to only one parameter, or a small number of parameters. We take it to be a mark of a poor causal representation of a system when, upon considering a near-variant of that system, we must change *many or all* of the parameters in our representation.

More formally, imagine that system *S* has two alternative parameterizations (such as the α/β and π/ϕ fitness structures). Call these sets of parameters **P** and **Q**. We assume these sets are of the same size, which is true for the case in question. A near-variant test can be applied as follows. Consider some small change to the system, yielding system *S**, and then consider the changes that must be made to each representation to accommodate this change and yield a dynamically sufficient representation of *S**. So **P** is replaced by **P***, and **Q** by **Q***. If, to achieve this, we must change *more* parameters in **P** than we do in **Q**, then parameterization **Q** is more natural, with respect to that near-variant.²

² An anonymous referee raised the possibility that some near-variants might be reflected in changes to the functional form of some equation(s) in a model, rather than the value of one or more parameters. We agree that this may occur, and it raises some more complicated possibilities. Suppose that two models **M**₁ and **M**₂ are representations of *S*, and a particular imagined modification to *S* can only be accommodated by changing the functional form of some equations in **M**₁ and **M**₂, but in different ways, yielding **M**₁* and **M**₂*. How do we determine which model has more naturally accommodated the imagined change? We agree that this possibility, if actualized, would raise a problem. We note, however, that often it will be possible to re-express the relation between **M**₁ and **M**₁* as one involving parameter values. For example, a linear relationship is a special case of many other functional relationships, with various parameters set to zero.

A thorough analysis of a pair of parameterizations will consider several near-variants, to see if one parameterization is superior across a wide range of alterations. If so, we view it as a superior representation of the causal structure of the original system. Essentially, we are combining the idea that a causally accurate representation is one in which individual parameters map to distinct aspects of the system being modeled, with the idea that imagined manipulations provide a way of assaying or testing for this property. We see this as one way of making more explicit the distinction between ‘mechanistic’ and ‘phenomenological’ models of a system—in cases where both models are inter-translatable. A model which maps parameters to distinct aspects of the system being modeled may be seen as faithful to the ‘mechanistic’ structure of the system, in one reasonable sense of that term.

The near-variant test assumes what may be a philosophically contentious view about the relationship between causation and localization. One way to defend such a view is to again work within an interventionist approach to causation. A cause is something that, when manipulated, yields changes to other variables. The notion of manipulation or intervention includes the idea that the factor in question be a reasonably localized feature of the system. Local features whose manipulation gives rise to significant changes in another variable are what John Campbell ([2007], [forthcoming]) calls ‘control variables’ for that effect. As Campbell argues, any macroscopic control variable will have some microphysical basis. That does not disqualify it from being a cause. What does tend to disqualify a macro-level feature, perhaps in favor of a lower level one, is a situation where what is referred to as a macro-level ‘factor’ is not reasonably localized but is highly distributed, or holistically realized, across many parts of the system. For Campbell, the role of localization in the concept of cause derives from the local nature of paradigm cases of interventions. Without endorsing an interventionist approach in general here, we do think that the interventionist view has described a genuine feature of causation and causal analysis.

It is important to understand what such a test is, and is not, supposed to do. Clearly the outcome of such a test will depend on what is taken to be a *near* variant of the system under consideration. Two people may disagree about the right kinds of perturbation to consider, and disagree so thoroughly that a ‘majority vote’ over the variants they consider relevant is either uninformative or impossible to apply fairly. The assessment of what counts as a near-variant may be influenced by causal assumptions about the system.³ For example, it

³ For example, the choice of Cartesian versus polar coordinates to describe the position of an organism in two dimensions could depend on how one supposes the organism in question moves. If movement occurs along cardinal directions, a Cartesian parameterization is more natural. If movement occurs as shifts around, toward, or away from the origin, a polar

may be affected by assumptions about which parts can change independently of each other. So we do not see this test as one that will extract causal judgments from an entirely acausal basis. Rather, the test *makes explicit* some elements of a largely tacit causal understanding of the system that we must have ahead of time. This understanding will be conditioned by experience with actual patterns of variation seen in the system, and also with mechanistic knowledge, but may draw on recognitional capacities that are hard to explicitly describe. Our test takes the deliverances of this informal causal understanding as raw material, makes them more precise, and reveals their further consequences.

We will illustrate the test with some simple cases. These are cases in which the correct description seems obvious ahead of time; in one case a breakdown of a collective into lower level entities seems entirely misleading, in the other it seems obligatory. Our aim is to use these easy cases to show how the near-variant test works, and then put it to work on more difficult ones. The first easy case we call the case of ‘pseudo-particles’. Consider a bacterial cell. This cell is comprised of a left half and a right half. Although it may seem peculiar to do so, let us think of these two halves as two ‘particles’ in a group that comprises the entire cell. If during a reproductive event the cell splits along its long-axis midpoint, then in a sense the left particle gives birth to two left-half particles, and the right particle gives birth to two right-half particles. (See Figure 1a.)

We can label the left half of the cell as A and the right half as B , and represent bacterial fission from either a collective (π/ϕ) or contextual (α/β) point of view (see Figure 1b). Thinking contextually, the fitnesses of the left half (or left pseudo-particle) and of the right half are both two ($\alpha_1 = \beta_1 = 2$). Thinking collectively, the original group of two particles gives rise to four particles, so $\pi_1 = 4$. Since left pseudo-particles comprise a half of every offspring cell produced, $\phi_1 = 1/2$. We then consider a near-variant (Figure 1c). Suppose the bacterial cell replicates at a faster rate, so that over the same period of time, four complete cells are produced. Thinking contextually, the left pseudo-particle then has four offspring ($\alpha_1 = 4$) and the right has four offspring ($\beta_1 = 4$). Thinking collectively, the number of offspring particles is now $\pi_1 = 8$, but the split of this productivity between left and right particles is still $\phi_1 = 1/2$. So to accommodate the near-variant, two parameters change in the contextual perspective, whereas only a single parameter changes in the collective perspective. At least with respect to this variant, the collective parameterization is more natural. This, of course, is what we would expect.

parameterization is more natural. Two different researchers could employ distinct ‘natural’ parameterizations because each has imagined (or witnessed) different ‘positional variants’.

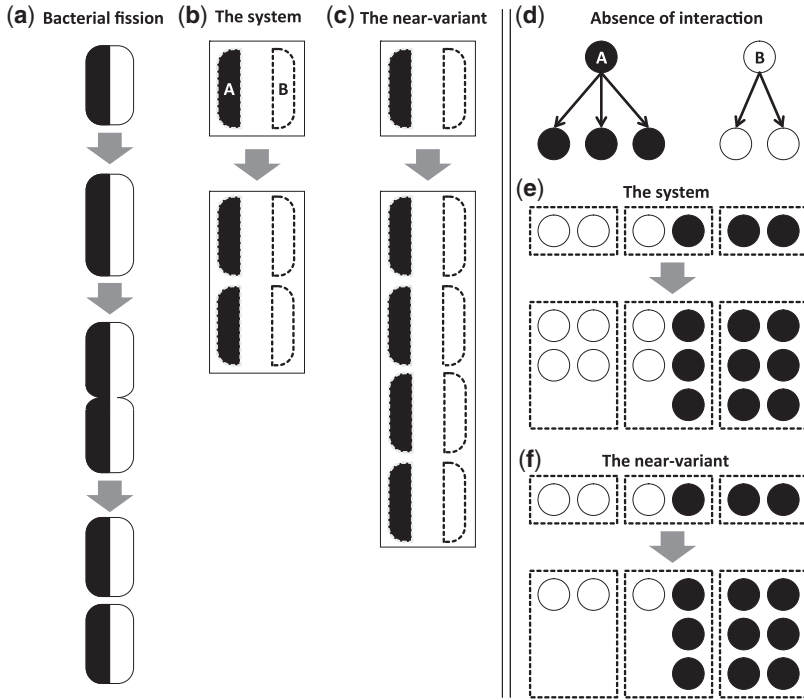


Figure 1. Pseudo-particles (a–c) and pseudo-groups (d–f).

The left and right half ‘particles’ here do not have any autonomy at all; there cannot be a group of two left halves, or a right half alone. Treating both halves as parts of a collective is far more natural.

This case can be contrasted with one in which the contextual parameterization is clearly superior; this is a case of ‘pseudo-groups’ (see Figure 1d; cases of this kind are also discussed in Sober [1984] and Nunney [1985]). Suppose that *A*-types always have three offspring and *B*-types always have two, regardless of their context. Individuals also tend to be found in the company of one other individual. There is no interaction between them, but we can, if desired, see every individual as part of a pair.

If we think contextually about this case, then $\alpha_i = 3$ for all relevant i , and $\beta_i = 2$ for all relevant i . Thinking collectively, the π_i parameters do vary according to i , because groups with more *A*’s are always more productive (specifically, $\pi_0 = 4$, $\pi_1 = 5$, and $\pi_2 = 6$). The sole non-trivial ϕ parameter, ϕ_1 , is $3/5$.

We then consider a near-variant: suppose each *B*-type produces only one offspring, not two. We will not work through the details, but it is easy to show that now the contextual parameterization is judged superior, as it can

accommodate the change with fewer modifications (two rather than three). So as earlier informal discussions had claimed regarding this case, the contextual parameterization (hence an individualist description of selection) is superior.

We now look at a case that is more substantial, an example of meiotic drive. In meiotic drive, one allele (A) gains an advantage over another (B) because the ‘driving’ allele is found in more than half of the gametes from heterozygote individuals (individuals of genotype AB). Such systems are generally modeled with genotype fitness differences and a drive parameter reflecting the advantage enjoyed by one allele in the heterozygote during gamete formation. Suppose in such a case we have the following genotype fitnesses: $W_{AA}=6$; $W_{AB}=5$; $W_{BB}=4$. The value of the driving allele d , reflecting the fraction of heterozygote matings in which A is the fertilizing allele, is $3/5$. Readers may note that this is essentially a multi-level parameterization; the genotype fitnesses are π values and the drive parameter is ϕ_1 . This is standard in population genetics (for discussion see Lloyd [2005]; Waters [2005]). One might now ask whether this orthodox representation of the meiotic drive case has a causal rationale. Note first that, as discussed earlier, it is possible to switch to a contextual parameterization, in which alleles are the only bearers of fitness. In this case, $\alpha_1 = \alpha_2 = 3$; $\beta_0 = \beta_1 = 2$. Note that this is the same fitness structure seen in the pseudo-group case (see Figure 1d–f, also Okasha [2004] for a similar example). The pseudo-group case seemed clearly to be one in which the contextual parameterization is more natural. This intuitive judgment was vindicated by a near-variant analysis. One might think that this meiotic drive case must be treated the same way as the pseudo-group case, as the fitness parameters are all identical. However, let us consider the near-variant test in the meiotic drive case. What would be a reasonable variant? One possibility is a change that affects the rate of drive and nothing else. This simultaneously changes α_1 and β_1 in the contextual perspective, but only ϕ_1 in the collective parameterization. This is because the model here features a simple zero-sum relation between the productivity of alleles in the heterozygote. We can compare this to a case where the advantage of one allele over another derives from superior viability in a competitive interaction between gametes after gamete formation, outside the father’s body, and in a scramble between gametes to fertilize the egg. Spawning in fish is an example; allele differences may then affect the viability of gametes in the water column. In that case, imagined perturbations to the viability of one gamete type (e.g. α_1) do not logically imply changes to the viability of the other (β_1). Consequently a contextual parameterization will be superior. So we see from consideration of the meiotic drive case that the fitness structure alone does not determine the most natural parameterization. Once we attend to the biological details, the near-variant test can sort instances of a single fitness structure into different causal categories.

We now apply this test to a case that has been the focus of much controversy, the case of competition between ‘altruists’ and ‘selfish’ individuals in a trait-group model. In a ‘trait group’ model, groups form by aggregation of lower level individuals at a certain stage in the life cycle, and the groups dissolve at another stage. When the composition of these groups affects the fitness of the lower level entities, some see this as a case of multi-level selection (Wilson [1980]; Sober and Wilson [1998]). Others see it as involving individual-level selection only (Maynard Smith [1976], [1998]), and others still have argued that it can be accurately described either way. In a model of this kind, suppose that all individuals have a ‘baseline’ absolute fitness of two offspring, but the presence of an altruist within a group allows other individuals in its group to have two extra offspring each. The altruist incurs a cost of one offspring. Assuming groups of $n = 3$, the consequences of a single selective episode are given on the left hand side of Figure 2.

We then consider a near-variant. Suppose that A individuals provide a benefit of *four* offspring to every other member of their group, at a cost of *two* offspring to themselves. This variant assumes that there is a linear relationship between cost and benefit. When we adjust the two parameterizations to accommodate such a change, we find they are *equally* sensitive. In both cases, five out of six parameters must be changed. The result is a tie.

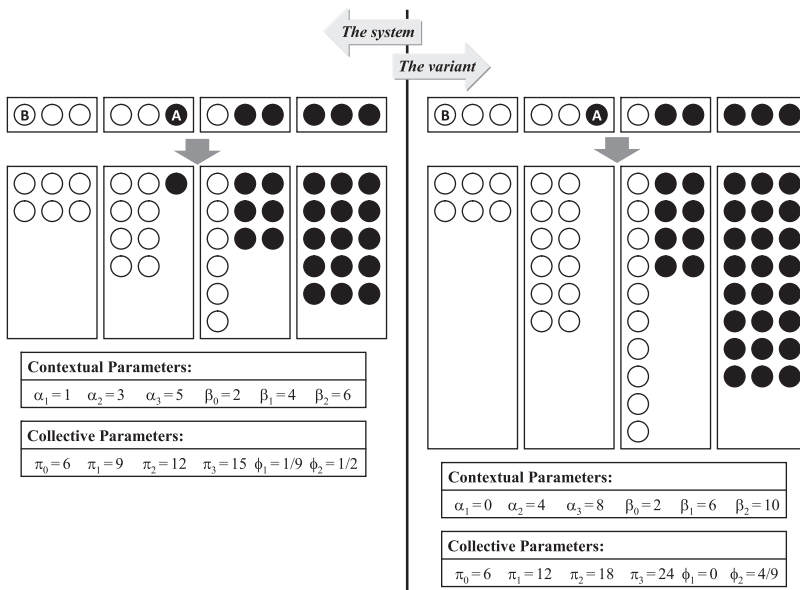


Figure 2. A trait-group model of altruism.

Not all near-variants to this system yield ties. If costs alone, or benefits alone, are altered, then the contextual parameterization fares better. And if costs and benefits are stipulated to be equal and group size is two, the collective parameterization wins. Given this mixture of results, and especially the tie seen in the linear case and some other cases (not shown), it seems that the near-variant test will not adjudicate between competing descriptions of a system of this kind.

Some might find this conclusion frustrating, but we think it is entirely appropriate. We are dealing with a case where individuals are not constitutively tied to their groups (as in the pseudo-particles case), and do have ‘a life of their own’. But they are also parties to an important form of interaction that unites the individuals within each group. The result is that the contextual and collective perspectives are found to be on at least roughly equal footing. This provides further support for the ‘pluralist’ interpretation of trait-group cases discussed above. This conclusion should be qualified by the recognition that in some cases of trait-group models, the biology of the system under discussion may make some particular class of near-variants relevant in a way that favors one parameterization or the other. The general features of the trait-group structure, however, do not favor one parameterization or the other. Further, the failure of the near-variant test in cases such as these may be informative in another way. When a near-variant test fails to discriminate alternative descriptions, and there is no missing empirical information which might change the verdict, it indicates a *partial* entanglement of the lower level entities into higher level collectives, with respect to their evolutionary role. This may be a sign that the system is on the road to an evolutionary transition.

4 State-variable Gestalt-Switching

The second main idea of this article is the possibility of a different kind of gestalt-switching. In the work discussed above, gestalt-switching involved two ways of looking at the *bearers of fitness*. Groups can be seen as fitness-bearers, or they can be seen as aspects of the context experienced by individuals. But the ‘multi-level’ description that results from recognizing group fitness in this way is one in which all the *accounting* is done in the currency of lower level particles. This applies to both the measurement of fitness and the representation of change. A fit group, in these models, is one whose constituent particles produce many new particles. In Damuth and Heisler’s terms, this yields an ‘MLS1’ model (Damuth and Heisler [1988]; Okasha [2006]). And, in the models above, the system is said to undergo evolutionary change when the frequencies of different kinds of particles change. There is no need to track the reproduction of groups as units in a model of this kind.

We now look at how one can make a switch to thinking of groups not only as fitness-bearers but also as the entities that make up the *population that is treated in Darwinian terms*—the population of entities in which there is variation, heredity, and differential reproduction. We see this as involving two moves: (i) representing groups in the *state-variables* of the model, and (ii) using fitness parameters that represent reproduction of groups by groups. Here we will focus mostly on the first of these.

The new kind of gestalt-switching we will call *state-variable gestalt-switching*, as opposed to the *parameter gestalt-switching* above. To see how the new kind of switch works, we must attend to a feature of the earlier models not discussed in detail so far. In a trait-group model of the type above, change is treated as a consequence of both a *fitness structure* and a *group frequency distribution*. This distribution is a set of $f_i(t)$ values, specifying the frequency of groups containing exactly i individuals of the A -type. So in these models, the frequencies of groups were not used to track change, but were one input into processes that yield change in particle frequencies. However, it would also be possible to write equations for change that give the new frequencies of each type of *group*, as a function of old frequencies of the groups and other parameters.

Schematically, if F_i is a function predicting change in the frequency of groups with i particles of the A -type, the equations would look like this:

$$\begin{aligned}
 f_0(t+1) &= F_0(f_0(t), f_1(t), \dots, f_n(t), \text{fitness parameters}) \\
 f_1(t+1) &= F_1(f_0(t), f_1(t), \dots, f_n(t), \text{fitness parameters}) \\
 &\dots \\
 f_i(t+1) &= F_i(f_0(t), f_1(t), \dots, f_n(t), \text{fitness parameters}) \\
 &\dots \\
 f_n(t+1) &= F_n(f_0(t), f_1(t), \dots, f_n(t), \text{fitness parameters})
 \end{aligned} \tag{1}$$

The frequency of each kind of group at the next time-step is a function of the frequencies of all the different kinds of groups at the previous time-step, along with parameters describing the fitnesses of entities in the system. We will work through an example, using a one-locus two-allele diploid population genetic model. In this case, the ‘particles’ are alleles (i.e. A and B), and group size is two. One way of writing such a model is to give equations for new allele frequencies as functions of old allele frequencies, fitness parameters, and a rule describing how alleles combine into genotypes. Assuming random mating and fair meiosis, the model yields equations of this kind:

$$\begin{aligned}
 \bar{W}p' &= W_{AA}p^2 + W_{AB}pq \\
 \bar{W}q' &= W_{BB}q^2 + W_{AB}pq
 \end{aligned} \tag{2}$$

Here W_{AA} , W_{AB} , and W_{BB} are absolute fitnesses of genotypes. The state variables, p and q , are the frequencies of particles. The primes signify frequencies in the following generation. However, the model could be reworked so that the frequencies of groups—diploid genotypes—become the state variables. Here x , y , and z are the frequencies of AA, AB, and BB groups, respectively.

$$\begin{aligned}\bar{W}x' &= W_{AA}(x^2 + xy + y^2/4) \\ \bar{W}y' &= W_{AB}(xy + 2xz + y^2/2 + yz) \\ \bar{W}z' &= W_{BB}(y^2/4 + yz + z^2)\end{aligned}\tag{3}$$

In Figure 3, we illustrate an example of this one-locus diploid population genetic life cycle (Figure 3a), where the focus is either on genes (Figure 3b) or genotypes (Figure 3c).

It will be possible, in the cases treated here so far, to use π and ϕ as the fitness parameters in these equations (and also possible to use α and β). But a full gestalt-switch will involve introducing a different kind of fitness parameter. We can use Ω_i as the number of offspring *groups* produced by a group with i A -types—this is an MLS2-type fitness parameter, in the sense of Damuth and Heisler ([1988]) (who also used ‘ Ω ’ to represent a group-level fitness parameter for MLS2). Depending on the case, we will also need a second set of fitness parameters, analogs to ϕ . The relations between π , ϕ , and Ω may be complex, depending on the mode of reproduction seen in a particular model.

So we have a different type of gestalt-switching, between models that track change in terms of particles and those that track change at a higher level. This state-variable gestalt-switching differs from parametric gestalt-switching in an important way. The α/β and π/ϕ parameterizations were fully interchangeable. In the case of our population genetic system, particle frequencies are given by group frequencies but not necessarily vice versa, and when the state variables are group frequencies the model has higher dimensionality. We can then reasonably ask how a model with group frequencies as state variables will earn its keep. When life cycles are dominated by the group phase, this representation may be more faithful to the biological details of the system, and when fitness depends on the frequencies of other groups, a higher level representation may be easier to formulate. An additional reason to consider state-variable gestalt-switching is to focus attention on issues of dynamic sufficiency. Models that are dynamically sufficient include information required to predict change over many time-steps.⁴ When considering state-variable gestalt-switching, the information needed for dynamic

⁴ The relation between dynamically sufficient models and abstract statistical summaries of change which are not dynamically sufficient is discussed in (Kerr and Godfrey-Smith [2002b]).

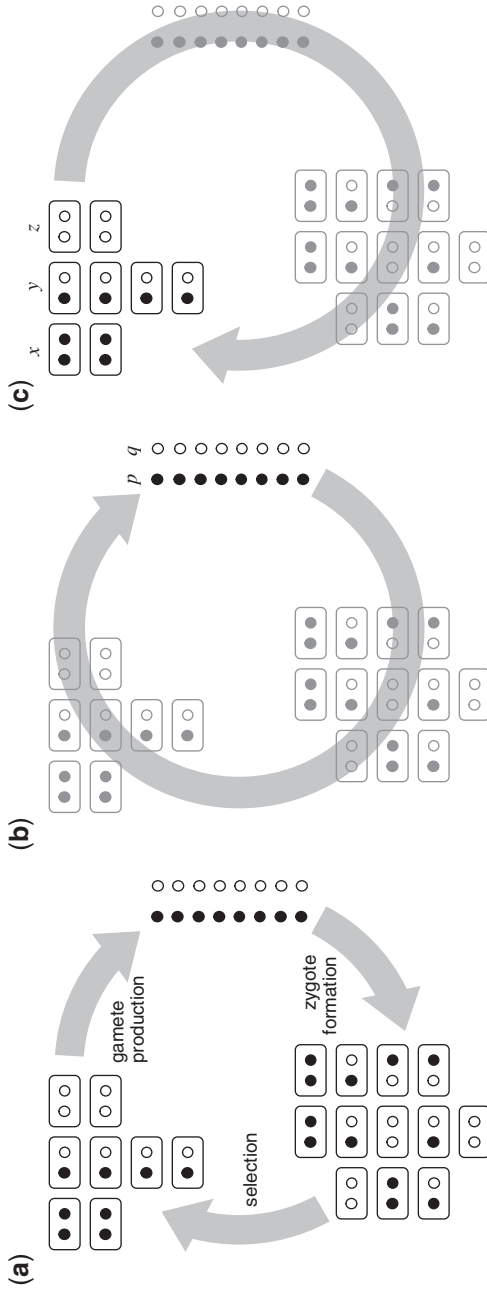


Figure 3. A diploid one-locus population genetic model, where black circles are A alleles and white circles are B alleles and rounded rectangles give diploid genotypes. (a) The life cycle involves production of gametes, formation of zygotes, and viability selection. (b) A model given by Equation (2) focuses on how the allele frequencies in one generation (p and q) yield the allele frequencies in the next generation (p' and q'). The state variables describe frequencies of *particles*. (c) A model given by Equation (3) focuses on how genotype frequencies in one generation (x , y , and z) yield the genotype frequencies in the next generation (x' , y' , and z'). Here, the state variables describe frequencies of *groups*.

sufficiency (i.e. how particles form groups and how groups generate particles) becomes apparent.

One role for these choices is to see them as alternative representations of a single case, highlighting different information. Okasha correctly notes ([2006], chapter 4) that the gestalt-switching discussed in our earlier work only concerns a multi-level description of the MLS1 kind; here we have made some steps towards the extension of the gestalt-switching framework to the richer sense of multi-level selection seen in MLS2 models.

This switch between frameworks also has a special relation to a particular kind of evolutionary process. Recent work has argued that the existence of a new level at which *bona fide* reproducing individuals are found is the mark of at least many evolutionary transitions (Buss [1987]; Michod [1999]). A transition often involves the appearance of a new 'Darwinian population' (Godfrey-Smith [2009]). A formal model might be seen as *recognizing* a particular Darwinian population when it tracks evolutionary change *by* tracking changes in the frequencies of types in that population, or by tracking change in the mean value of a characteristic of members of that population.

We suggest that exploring the relations between models with different state-variables might be useful when dealing with systems that have significance for evolutionary transitions. One such case is *Volvox carteri* (Michod et al. [2003], Michod [2005]). This organism is a colonial green algae that lives in ponds and lakes as hollow balls of cells. Each colony may contain daughter colonies, and even grand-daughter colonies. Daughter colonies can form by asexual reproduction. The colony itself is made up of 'somatic' and 'generative' cells. A generative cell undergoes several mitotic divisions to produce the daughter colony, which invaginates and later inverts inside the parent. Daughters are released when the parent colony dissolves. In this system we see a complex relationship between cell- and colony-level evolutionary processes, and we suggest that one strategy for understanding it would be to employ gestalt-switching with respect to both fitness parameters and state variables. Offspring groups are born as wholes when the parent dies, so there is visible group-level reproduction here, despite the simplicity of the group-level structures. Another empirical system that suggests itself for this treatment is the slime mold *Dictostylium discoideum* (Buss [1987]; Strassmann and Queller [2007]).

5 Representing an Evolutionary Transition

We conclude this part of the discussion by illustrating the roles that might be played by the different parameters and modeling tools described in this article in the representation of an evolutionary transition. The relations between stages of the transition and the tools discussed are summarized in Table 2.

Table 2. Transition stages and accompanying models

Stage	Biology	Model features
Stage 1	Interactions among particles without groups	Particle state variables, α/β fitness parameters
Stage 2	Interactions in groups	Particle state variables, α/β or π/ϕ fitnesses
Stage 3	More cohesive groups	Particle state variables, π/ϕ fitnesses via NV test?
Stage 4	Groups function as biological units, especially in reproduction	Group-level state variables, π/ϕ fitnesses
Stage 5	Groups function as biological units, especially in reproduction	Group-level state variables, Ω fitnesses (plus others as needed).

Imagine an initial state (Stage 1) with a population comprising A and B particles that interact with neighbours, but without forming discrete groups. Particle frequencies are used as state variables in a model of such a system, and particle fitnesses are represented with α and β parameters. (There is not even a formal possibility of using the parameters π and ϕ : Maynard Smith [2002]; Godfrey-Smith [2008]). But suppose that interactions become cooperative, and also organized into bounded groups (Stage 2). We can now use both π/ϕ and α/β parameters.

Once the groups become cohesive, a near-variant test may favor a π/ϕ representation (Stage 3). Then, as groups come to function more and more as units in their own right, especially in reproduction, a switch in *state-variables* may be motivated (Stage 4). We now describe evolution as change in a population of groups. But fitnesses are still expressed in terms of particles. So it may be sensible or (depending on the case) necessary to switch to using Ω fitness parameters. We are now tracking change in a population of groups, and we explain change in terms of the differential rates of reproduction *of* groups *by* groups (Stage 5). Thus, as in (Michod [2005]) and (Okasha [2005]), the relation between MLS1 and MLS2 models is understood in temporal, before-and-after terms.

Then suppose that the independent evolutionary role of the original particles is entirely suppressed. Groups ‘breed true’ when they produce new groups, and do so asexually. Then we can treat groups as integrated entities whose composition with respect to the original ‘particles’ need no longer be tracked. We have come full circle; the groups now behave like a new set of particles. If we then become interested in ways in which these higher level units interact with *each other*, we would have reason to re-introduce the original α_i and β_i parameterization at the higher level.

A key difference between this schematic pattern and actual-world cases is that integrated collectives in the actual world tend to engage in *sexual* reproduction. That complicates the relation between fitness parameters, and prevents the simple return to an α/β representation imagined here as our final stage. Actual-world transitions tend not to ‘come full circle’ in the way our schematic one does.

6 Conclusion

We have discussed two tools which have general utility when modeling levels of selection problems, and a special relation to evolutionary transitions. The near-variant test can be used to distinguish the causal accuracy of models even when they are predictively equivalent. Some controversial cases cannot be easily decided by the near-variant test, and this may furnish part of an explanation of *why* people have different intuitions about these cases. The near-variant test may also cast light on why particular modeling traditions gravitate toward particular parameterizations—population genetics embracing a collective treatment of groups of alleles; evolutionary game theory opting for a contextual approach to interacting pairs. In population-genetic models, most biologically plausible variations have consequences for all members of a group. This is not so with game-theoretic models. State-variable gestalt-switching makes possible a different kind of dual perspective on an evolving system. By noting the relations between models using different state-variables, as well as different fitness parameters, we can better understand what is involved in recognizing higher levels in the evolutionary process, and better understand the transitions by which higher level individuals and populations arise.

Acknowledgements

We are grateful to all the participants at the symposium at PSA 2006 on ‘Organisms Versus Models in the History of Life’, and to three anonymous referees, for helpful comments.

Peter Godfrey-Smith
Department of Philosophy
Harvard University
Cambridge, MA 02138, USA
pgs@fas.harvard.edu

Benjamin Kerr
Department of Biology,
University of Washington
Seattle, WA 98195, USA
kerrb@uw.edu

References

- Buss, L. [1987]: *The Evolution of Individuality*, Princeton: Princeton University Press.
- Campbell, J. [2007]: 'An Interventionist Approach to Causation in Psychology', in A. Gopnik and L. Schulz (eds), *Causal Learning: Psychology, Philosophy and Computation*, Oxford: Oxford University Press, pp. 58–66.
- Campbell, J. [forthcoming]: 'Causation in the Mind', presented as Alfred North Whitehead Lectures, Harvard University, 2009.
- Dugatkin, L. and Reeve, H. K. [1994]: 'Behavioral Ecology and Levels of Selection: Dissolving the Group Selection Controversy', *Advances in the Study of Behavior*, **23**, pp. 101–33.
- Damuth, J. and Heisler, L. [1988]: 'Different Formulations of Multi-Level Selection', *Biology and Philosophy*, **3**, pp. 407–30.
- Godfrey-Smith, P. [2008]: 'Varieties of Population Structure and the Levels of Selection', *British Journal for the Philosophy of Science*, **59**, pp. 25–50.
- Godfrey-Smith, P. [2009]: *Darwinian Populations and Natural Selection*, Oxford: Oxford University Press.
- Kerr, B. and Godfrey-Smith, P. [2002a]: 'Individualist and Multi-Level Perspectives on Selection in Structured Populations', *Biology and Philosophy*, **17**, pp. 477–517.
- Kerr, B. and Godfrey-Smith, P. [2002b]: 'On Price's Equation and Average Fitness', *Biology and Philosophy*, **17**, pp. 551–65.
- Lloyd, E. A. [2005]: 'Why the Gene Will Not Return', *Philosophy of Science*, **72**, pp. 287–310.
- Lewis, D. K. [2000]: 'Causation as Influence', *Journal of Philosophy*, **97**, pp. 182–98.
- Maynard Smith, J. [1976]: 'Group Selection', *Quarterly Review of Biology*, **51**, pp. 277–83.
- Maynard Smith, J. [1998]: 'The Origin of Altruism', *Nature*, **393**, pp. 639–40.
- Maynard Smith, J. [2002]: 'Commentary on Kerr and Godfrey-Smith', *Biology and Philosophy*, **17**, pp. 523–7.
- Maynard Smith, J. and Szathmáry, E. [1995]: *The Major Transitions in Evolution*, Oxford: Oxford University Press.
- Michod, R. [1999]: *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*, Princeton: Princeton University Press.
- Michod, R. [2005]: 'On the Transfer of Fitness from the Cell to the Multicellular Organism', *Biology and Philosophy*, **20**, pp. 967–87.
- Michod, R., Nedelcu, A. and Roze, D. [2003]: 'Cooperation and Conflict in the Evolution of Individuality IV. Conflict Mediation and Evolvability in *Volvox carteri*', *BioSystems*, **69**, pp. 95–114.
- Nunney, L. [1985]: 'Group Selection, Altruism, and Structured-Deme Models', *American Naturalist*, **126**, pp. 212–30.
- Okasha, S. [2004]: 'Multi-Level Selection, Covariance and Contextual Analysis', *British Journal for the Philosophy of Science*, **55**, pp. 481–504.
- Okasha, S. [2005]: 'Multi-Level Selection and the Major Transitions in Evolution', *Philosophy of Science*, **72**, pp. 1013–25.
- Okasha, S. [2006]: *The Levels of Selection Debate*, Oxford: Oxford University Press.

- Pearl, J. [2000]: *Causality*, Cambridge: Cambridge University Press.
- Sober, E. [1984]: *The Nature of Selection*, Cambridge MA: MIT Press.
- Sober, E. and Wilson, D. S. [1998]: *Unto Others: The Evolution and Psychology of Unselfish Behavior*, Cambridge MA: Harvard University Press.
- Sterelny, K. [1996]: 'The Return of the Group', *Philosophy of Science*, **63**, pp. 562–84.
- Strassmann, J. E. and Queller, D. C. [2007]: 'Altruism among Amoebas', *Natural History*, **116**, pp. 24–9.
- Waters, C. K. [2005]: 'Why Genic and Multi-Level Theories are Here to Stay', *Philosophy of Science*, **72**, pp. 311–33.
- Wilson, D. S. [1980]: *The Natural Selection of Populations and Communities*, Menlo Park: Benjamin/Cummins.
- Woodward, J. [2003]: *Making Things Happen*, Oxford: Oxford University Press.