Dynamics, Symmetry, and the Levels of Selection

Benjamin C. Jantzen
Virginia Tech

Abstract

Most attempts to answer the question of whether populations of groups can undergo natural selection focus on properties of the groups themselves rather than the dynamics of the population. Those approaches to group selection that do emphasize dynamics lack an account of the relevant notion of equivalent dynamics. I present a new approach to identifying instances of evolution by natural selection that is based upon dynamical symmetries. I apply the symmetry method to arrive at an affirmative but qualified answer to the question of group selection.

1. Introduction

Suppose it's possible for a population of one type of biological entity to evolve by natural selection. Is it also possible for a population of groups of those entities to evolve as well? More broadly, what populations of biological entities can be said to undergo a process of evolution by natural selection? Populations of organisms? Of alleles? Of demes? These questions concern the problem of the 'levels of selection' (Okasha, 2006, Sober and Wilson, 1994, Sterelny and Kitcher, 1988, Wilson and Sober, 1989), and in this essay, I argue for a new approach to addressing this problem. In particular, I claim that we should recognize questions about group selection as questions about types of dynamical process, and argue that dynamical processes are best classified on the basis of their symmetries. Once we specify the symmetries of natural selection, it follows that in certain circumstances, populations of groups can in fact instantiate such a process.

To provide some intuitive motivation for the dynamical symmetries approach, consider the following toy example. Imagine I have an agar plate with multiple bacterial colonies on it, and that I wish to know how many kinds of bacterial colonies there are. One thing I might do is take a sample of each colony and use it to inoculate a fresh plate—one for each colony—that is laced with penicillin. I could then examine the plates after a while to see how quickly each colony grew. If some died and the others grew at more or less the same rate as before, I would say that I had at least two kinds on the original plate. Similar manipulations are possible with other antibiotics, but less destructive techniques are available as well. I could, for instance, vary the concentration of nutrients in the growth medium and sort colonies on the basis of whether or not their growth rate is affected by the variation. With a little more effort, I could also vary nutrient

concentrations while keeping either the differences or relative proportions fixed. In each case, I could classify colonies on the basis of whether or not the rate of growth was affected by the change in medium. A variety of more or less complex patterns are possible. While a colony's growth rate may be changed by intervention on lactose concentration alone, it might be unaffected if lactose and glucose are varied together.

The kinds determined in this way are dynamical kinds. What my experiments discerned are not kinds of static colonies but rather kinds of metabolic processes. This is clear if we focus attention on the colonies whose metabolic processes remained constant under intervention. What I've discovered is that for those colonies, their metabolism—however complex it may be—possesses a dynamical symmetry. The term 'dynamical symmetry' denotes any transformation of a system that leaves unchanged the way in which the states of that system unfold through time. In this case, it is the unfolding of metabolic processes through time that remains unchanged under a transformation of nutrient concentration.

If the growth of a colony exhibits a bunch of dynamical symmetries, S, we'll call that type of colony an S-metabolizer. With this terminology and the notion of a dynamical symmetry we can provide answers to questions of the following sort: Is the combination of two S-metabolizers also an S-metabolizer? In this case, the answer is obviously yes—if the growth rate of each colony is independent of the transformations in S, then so is the growth rate of the two combined. The approach is quite general; symmetries can tell us more about colony dynamics. For instance, we can also immediately see that the no combination of an S-metabolizer and a non-S-metabolizer can be an S-metabolizer. While in the context of my toy example such questions are trivial to

answer, they take on a greater significance when we recognize that asking about metabolizers is closely analogous to asking about evolving populations. In particular, these examples illustrate the utility and precision of appealing to symmetries to discriminate dynamical types. More to the point, they suggest a precise and general method for identifying instances of a particular kind of process. Below, I flesh out a symmetry approach to the question of the levels of selection, a question which we can recast as follows: How can we carve up the biological world into populations of entities such that the dynamics of change in those populations respects the symmetries characteristic of evolution by natural selection?

To answer this revised question, I begin by looking at a previous approach to the levels of selection that puts dynamics up front, and identify where this approach goes awry. I then introduce a bit of formal machinery that will allow me to propose a characterization of evolution by natural selection in terms of dynamical symmetries. Finally, I attempt to answer the question of levels of selection, and consider a couple of examples for which group selection either does or does not take place.

2. How not to take dynamics seriously

One notable attempt to use dynamics to address the levels of selection is that of Godfrey-Smith and Kerr (2002) who develop a framework for talking about processes at both the 'individual' and 'group' level. Whatever the merits of this framework, however, it does not in fact provide the resources to answer the question about group selection with which we are concerned.

Understanding why this is the case will go a long way toward motivating the dynamical symmetry approach, and so it is worth considering the framework in some detail.

To begin with, Godfrey-Smith and Kerr (hereafter GSK) ask us to consider an infinite population of biological entities that are of two types, A and B. These entities, whatever they may be, we'll call 'particles' as GSK do in order to avoid confusing use of terms like 'individual' or 'organism'. These particles are supposed to aggregate in groups, all of fixed size n, where each group type is determined by the number i of type A particles that belong to it. The frequency of groups of type i is given by f_i. The particles and groups reproduce in discrete generations—groups dissolve, particles replicate in a single, panmictic mass, and then reform groups. According to GSK, we have two choices when modeling the behavior of this system over time. On the one hand, we can assign context-dependent fitnesses to each particle on the basis of its type and the type of group in which it occurs. Thus, α_i is the fitness of an A particle in a group with i A's and β_i is the fitness of a B particle in a group with i B's. In this approach then, groups are treated as local environments.

Let p stand for the frequency of A particles in the entire meta-population and q for the overall frequency of B particles. Then, GSK tell us, the dynamics of the p and q are captured by the following recurrence relations:

$$p(t+1) = \frac{1}{\overline{w}} \sum_{i=1}^{n} f_i(t) \frac{i}{n} \alpha_i \tag{2.1}$$

$$q(t+1) = \frac{1}{\overline{w}} \sum_{i=0}^{n-1} f_i(t) \frac{(n-i)}{n} \beta_i \tag{2.2}$$

The variable $\bar{w}$ is shorthand for the average particle fitness which is given by

$$\bar{w} = \sum_{i=1}^{n} f_i(t)\frac{i}{n}\alpha_i + \sum_{i=0}^{n-1} f_i(t)\frac{(n-i)}{n}\beta_i \qquad (2.3)$$

Equations (2.1) and (2.2) are not dynamically sufficient unless we specify the functions $f_i(t)$. In the cases considered by GSK, the $f_i(t)$ are assumed to be functions of $p(t)$.

In constructing (2.1) and (2.2), we took the perspective of the particles and treated groups as environments over which particle fitness varies. Of course, if we're going to be able to entertain the question of levels of selection, we're going to have to have a way to adopt the group perspective as well. GSK claim that we can do so in the following manner. To each group we assign two parameters:

$\pi_i =$ total number of copies from a group with i A types

$\varphi_i =$ (number of A copies in a group with i A types) / (total number of copies in a group with i A types)

The parameter π_i, represents a "group productivity" which is purportedly analogous to particle fitness. The latter parameter φ_i is the fraction of particles of type A produced by a group with i A-types. It has no analog from the particle perspective. Unsurprisingly, these parameters can each be written in terms of α_i and β_i, though the exact expressions are immaterial here. Using the new parameters as defined (and again assuming the functions f_i are known), we can write down a new set of dynamical equations for groups:

$$\bar{\pi}p(t+1) = \sum_{i=1}^{n} f_i(t)\phi_i\pi_i \qquad (2.4)$$

7

$$\bar{\pi}q(t+1) = \sum_{i=0}^{n-1} f_i(t)(1-\phi_i)\pi_i \qquad (2.5)$$

In this case, $\bar{\pi} = \sum_{i=0}^{n} f_i(t)\pi_i$.

As I suggested above, this setup does not allow us to answer the question of levels of selection as it was framed in Section 1. While it is true that Equations (2.4) and (2.5) represent some facts about groups, a glance at the definition of φ_i indicates that we are still firmly rooted in the particle perspective. Groups are acknowledged as "fitness structure" in the population, but this structure is ultimately still described with reference to particles. In fact, Equations (2.4) and (2.5) only explicitly track particle frequencies. What matters for the question of levels of selection as I've posed it is whether or not the dynamics of the population of groups is a process of Darwinian evolution when those dynamics are described exclusively in terms of a population of groups. To answer this question, these equations need a little rearranging.

To begin with, it is essential to note that the f_i are the frequencies of group types. If we want to know how the population of groups changes over time, we need to keep track of the f_i, not p and q. So let's suppose that the population of particles evolves in time according to Equations (2.1) and (2.2). Recall that in order for those equations to be dynamically sufficient, we have to know what the frequencies of groups are in terms of the frequencies of particles. Let's assume that groups are formed at random in the new generation following dissolution of the groups of the preceding generation. On this assumption, we have:

$$f_i(t) = \binom{n}{i}(p(t))^i(1-p(t))^{n-i} \qquad (2.6)$$

8

For simplicity, I will assume that $n = 2$. That is, each group is composed of just two particles. From Equations (2.1) – (2.5) we can then find explicit expressions for the dynamics of the group frequencies:

$$f_0(t+1) = \frac{\left(2\beta_0 f_0(t) + \beta_1 f_1(t)\right)^2}{\left(2\left(\alpha_2 f_2(t) + \beta_0 f_0(t)\right) + \left(\alpha_1 + \beta_1\right) f_1(t)\right)^2} \tag{2.7}$$

$$f_1(t+1) = \frac{2\left(\alpha_1 f_1(t) + 2\alpha_2 f_2(t)\right)\left(2\beta_0 f_0(t) + \beta_1 f_1(t)\right)}{\left(2\left(\alpha_2 f_2(t) + \beta_0 f_0(t)\right) + \left(\alpha_1 + \beta_1\right) f_1(t)\right)^2} \tag{2.8}$$

$$f_2(t+1) = \frac{\left(\alpha_1 f_1(t) + 2\alpha_2 f_2(t)\right)^2}{\left(2\left(\alpha_2 f_2(t) + \beta_0 f_0(t)\right) + \left(\alpha_1 + \beta_1\right) f_1(t)\right)^2} \tag{2.9}$$

To reiterate, what (2.7) – (2.9) represent are the dynamics of a population of groups as it changes in time, assuming that the particles composing the groups are governed by (2.1) and (2.2). These expressions, not Equations (2.4) and (2.5), are the analogues of the equations governing particle dynamics. Note that the population of groups is presumed to occupy a single selective environment—there is only one group context. We should thus expect there to be a single group fitness for each type of group. In fact, we'll have reason later (in Section 4 below) to identify group fitnesses with the coefficients of the frequencies appearing in the denominators of each expression on the right-hand side, namely α_2, β_0, and $(\alpha_1 + \beta_1)$. These happen to correspond with the parameters π_i introduced as 'group fitnesses' by GSK. But this cannot be determined by simple inspection of the equations. Only the group frequencies are unambiguously recognizable. What's more, the form of, say Equation (2.7) is quite different from that of Equations (2.1) or (2.2). In particular, frequencies appear only to first order on the right hand side of (2.1) and (2.2). This is not the case for any of the equations (2.7) – (2.9). Whatever parameter combination we

9

pick for the group fitnesses, there is no getting around the fact that the dynamics governing change in the population of groups looks different from the dynamics of particles.

To take stock so far, we began by assuming that particles comprise an evolving population and stipulated the dynamics governing this evolution. With the additional assumption of a particular connection between particles and groups, we were then able to produce a model of the dynamics of the population of groups in terms of group properties alone—properties that could be assessed without any knowledge that the groups consist of particles. It remains to be seen whether this population of groups can be said to evolve via natural selection and thus constitute what I'll call—in Godfrey-Smith's (Godfrey-Smith, 2009) felicitous idiom—a Darwinian population. More specifically, the question we have yet to address is whether the dynamics expressed by Equations (2.7) – (2.9) is of the right sort to count as evolution by natural selection. Note that I am not asking about whether the population possesses the right sort of properties apart from its dynamics—it is not directly relevant to our question whether, for instance, there is a mechanism for generating heritable variation. Of course, if it turns out to be a Darwinian population, then we have good reason to suspect that there is such a mechanism. But knowing there is such a mechanism is insufficient to guarantee any particular dynamics, and it is the dynamics that makes the population Darwinian.

Though relatively few have done so, I am hardly the first to place this kind of emphasis on dynamics. For instance, Wagner and Laubichler (2000, 32) identify evolutionary characters (or traits) as follows: "given the equations that describe the dynamics of natural selection among individuals…how can we lump the organism 'types' (be it genotypes or phenotypes) into

equivalence classes (i.e. character states), such that the dynamics of these abstract types is still predicted by the same equations without any loss of generality." But what does it mean for the states of two different populations to be predicted by the same equations without loss of generality? Presumably, natural selection encompasses more than evolution in a haploid population with two alleles and n fitness environments as described by Equations (2.1) and (2.2) above. But even the addition of a third allele or just one more fitness environment would change these equations. Obviously, the 'sameness' of dynamical equations cannot be taken too literally here. But how can we decide whether or not the dynamics of two changing populations is the 'same' in some more salient sense? A satisfying answer to this question is provided by the formal notion of symmetry structure.

3. Dynamics, symmetry, and individuality

Suppose we are interested in Newtonian gravity instead of Darwinian evolution. The analog to our question about levels of selection is the following: Which collections of objects comprise a gravitating system? To see the analogy more closely, we might ask whether collections of gravitating systems can also compose a gravitating system. One way we might approach the problem is similar to the way in which we used Equations (2.7) – (2.9) above. That is, we can begin with a gravitating system of, say, n massive particles with dynamics described by Newton's laws of motion in conjunction with the inverse square law of gravitation. As with our biological 'particles', we can then redescribe the system in terms of m groups of particles (it doesn't matter whether we assume the same number of particles in each group or not). From the original dynamical equations, we can then deduce new equations describing the evolution of the

meta-system entirely in terms of groups. I won't present an explicit calculation here, but the upshot is that we will be confronted by the same sort of problem we faced above: in what specific ways and to what degree must the dynamical descriptions be similar for us to declare them both gravitating systems.

To answer this puzzle we should focus not on the details of the dynamics, but rather on the dynamical symmetries. Simply put, a *dynamical symmetry* of a system is any physical transformation of the system that is invisible to the dynamics. To put it a little more precisely, a transformation, σ, is a symmetry if it makes no difference whether we apply σ to the system and then advance the state according to the dynamical laws, Λ, or if we instead apply the laws, Λ, to advance the state and then transform the result with σ. Either way we end up with the system in the same state.[1] In classical mechanics, for instance, displacement by a constant distance in space is a dynamical symmetry. Imagine a system that exhibits a simple ballistic trajectory, such as a cannonball fired from a cannon. We can describe the initial state of the system with the coordinates of a single point and an initial velocity. Advancing the system using the dynamical laws yields a series of states that trace out a parabola in space. This is depicted by the two plots on the left-hand side of Figure 1. In the upper left, the initial position of the cannonball is indicated by a circled point. The result of advancing this state through time with the dynamical laws is shown in the plot below it. Now suppose that we transformed the original system by shifting the cannon to its right by 10 m. This transformation, σ, of the initial state is indicated by the plot in the upper right corner of Figure 1. If we then advance the state through time, we would see the ball follow a parabolic trajectory as shown in the lower right corner. But each point on this trajectory represents exactly the state we would have gotten were we to have

---

[1] This is similar to the notion of symmetry developed in (Rosen, 1995).

transformed the corresponding point on the original trajectory by σ. That is, it doesn't matter whether we move the cannon and fire the cannonball, or fire the cannonball and then move it. Either way, we get the same state. In this sense, the dynamics of motion are insensitive to spatial translations. Wherever you are in the world, the path of a ballistic cannonball looks the same.
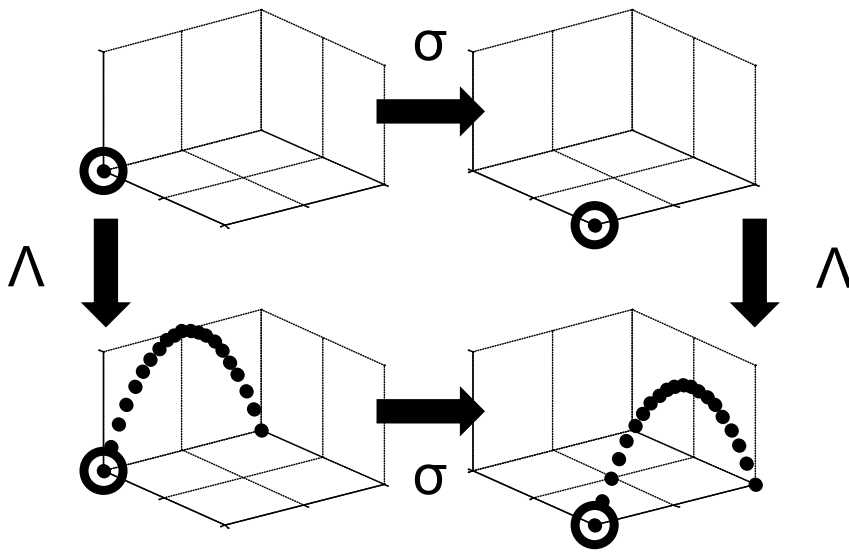

Figure 1.

Consider instead the bacterial example with which I began. Each colony is a dynamical system. While lots of properties of the colony change with time, let's focus strictly on biomass (usually approximated by measuring area). As time progresses, suppose that a given colony on a particular growth medium adds biomass at a constant rate. If we change the growth medium—by adding antibiotic, for example—this constitutes a transformation of sorts. If the growth rate of the colony is insensitive to the concentration of the antibiotic, then it will not matter whether I change the antibiotic concentration of the medium and then grow the colony, or grow the colony

and afterward add the antibiotic. I get the same state either way, and so this transformation is a symmetry of the colony dynamics.

Dynamical symmetries weave together into what I call symmetry structures. Suppose that σ_1 and σ_2 are symmetries of the dynamics of interest. Then so is $\sigma_2 \bullet \sigma_1$, where this notation denotes the operation obtained by first applying σ_1 and then σ_2. In fact, it is easy to see that this generalizes to any composite of symmetry transformations—the symmetries of a set of dynamical equations form a group in the mathematical sense.[2] A particular group structure along with a list of the specific transformations that compose it constitutes the symmetry structure of a dynamics. A concrete biological system respects the symmetry structure of a dynamics if its states preserve this structure under the indicated transformations.

Symmetry structures offer a way of classifying individual processes according to a dynamics of interest. For a particular system to instantiate a particular dynamics, it is a necessary condition that the system manifests the symmetry structure associated with that sort of dynamics. So, for instance, if a system is to be a Newtonian gravitating system, then it must manifest the dynamical symmetry structure of Newton's laws of motion and gravitation. In particular, the dynamics must be unaffected by rigid translations and rotations of the system (and a handful of other such transformations) as well as combinations of these transformations. Prima facie this might sound like a circumspect way of saying that two systems have the same dynamics if they "obey the same equations." That would be deeply unsatisfying since we already know that the systems we want to lump together as Darwinian populations do not all obey exactly the same dynamical

---

[2] There is no reason a priori to rule out dynamical symmetries that have no inverse, and so it would be more accurate to say that symmetries form semigroups.

equations. But the appeal to symmetry does not merely group systems by their specific dynamical equations. In fact, one cannot do so on the basis of symmetry conditions alone since the fact that a system respects a particular symmetry structure is not always a sufficient condition for that system to be governed by a particular dynamics—two systems that respect the same symmetry structure might differ in dynamical details. This flexibility is appealing. For one thing, there are no difficulties under the symmetry approach with accommodating different numbers of variables corresponding to, e.g., different numbers of alleles. This is because symmetry structures are abstract and may be instantiated by all kinds of concrete systems provided that the transformations are appropriately specified. In the gravitating system example, it doesn't matter how many gravitating bodies we consider—the notion of a rigid translation can be precisely specified with a schema such as: "add a constant vector to every position." However, the fact that symmetry structures impose a non-trivial equivalence relation on dynamical systems means that to insist on identifying dynamical kinds on the basis of symmetry is to make a substantive claim that requires defense. Why classify dynamical processes this way and not some other way?

The reason we should take the equivalence relations imposed by symmetry structures seriously as dynamical kinds is that each such kind picks out a class of units within a domain of inquiry that are independent of one another. That is, if we assume that transformations can be restricted to a particular system of interest, it is the case that if a system respects the symmetry structure of a given dynamics then its states cannot depend upon the states of other such systems. So for example, the state of any Newtonian gravitating system as it evolves through time must be independent of the state of any other gravitating system. Conversely, if the state of one collection of massive bodies influences the state of another, neither collection is a Newtonian gravitating

system. To return to my agar plate example, any partition of colonies into S-metabolizers separates colonies into units that grow at rates independent of one another. If the growth rates of two colonies are tied together, then neither is an S-metabolizer.

This way of dividing systems thus reflects a genuine division in the world. In an intuitive sense then, each process picked out as an instance of a dynamical kind is an individual in the sense of unit that bears properties independent of others. Such individuals are precisely the sorts of entities over which one can quantify when stating (or seeking) higher-level empirical regularities. Furthermore, it is lot easier to determine whether or not a system respects a symmetry structure than it is to determine the dynamical details of the system in terms of its internal variables. For instance, it is quite easy to identify an S-metabolizer via a few experiments with agar plates. It is very hard to determine precisely how the overall metabolic rate of the population of bacteria is caused by the distribution and concentration of nutrients. Thus, not only does the symmetry approach pick out plausible natural kinds, but it is easier to learn these kinds than it is to learn full dynamics.

4. Evolutionary individuals

So what is a Darwinian population? I have argued that the question of the levels of selection can be settled on symmetry grounds: either the dynamics of a population described at some level of the biological hierarchy respects the symmetries of evolution by natural selection—in which case selection does occur at that level—or it violates one or more of these symmetries, in which case selection cannot be said to operate at that level. To ascertain whether or not a particular

population is a Darwinian population—a population undergoing a process of natural selection—we need only ascertain whether the dynamics of the population obeys the symmetries of evolution by natural selection. Of course, this means addressing the daunting question of what the 'fundamental' dynamics of evolution by natural selection are.

Restricting our attention to the deterministic models of population dynamics (which generally assume infinite populations), the list of fundamental equations is topped by the so-called 'Replicator-Mutator Equation' or RME. This equation comes in different forms for haploid and diploid populations:[3]

The RME for discrete generations (haploid):

$$f_i(t+1) = \frac{1}{\bar{w}} \sum_j q_{ij} w_j f_j(t) \tag{4.1}$$

The RME for discrete generations (diploid):

$$f_i(t+1) = \frac{1}{\bar{w}} \sum_{jk} f_j(t) f_k(t) F_{jk} Q_{jki} \tag{4.2}$$

The variable f_i refers to the frequency of a type (such as a genotype) and w_i refers to the fitness of that type. The term q_{ij} is the probability that an individual of type i produces a unit of type j in haploid reproduction, while Q_{jki} is the probability that a mating of type j with type k individuals produces an offspring of type i. In Equation (4.2), $F_{jk} = F_{kj}$ is the expected number of offspring

---

[3] The haploid version of the RME is taken from (Page and Nowak, 2002, 97). The diploid version was derived by the author by using the discrete version of terms appearing in the continuous time RME for sexual reproduction as presented in Appendix A of the same paper.

any pair of individuals of types j and k produce together, and so $w_i = \sum_j f_j F_{ji}$. In both cases,

$$\bar{w} = \sum_i w_i f_i.$$

Most if not all of the standard models in population genetics can be cast as special cases of either Equations (4.1) or (4.2). In fact, the haploid RME can be treated as a special case of the diploid RME for which $\sum_k x_k F_{jk} Q_{jki} = f_j q_{ji}$. The Price Equation, which describes the change in mean character value across a single generation, has been touted for its generality (Godfrey-Smith, 2009, Okasha, 2006). But this equation too can be derived from Equation (4.1) (Page and Nowak, 2002). All of these models therefore inherit the symmetry structure of the RME, and so there is good reason to take this symmetry structure as characterizing the process of evolution by natural selection. Taking the RME to be fundamental leads to the following:

Symmetry Condition: For a population to be a Darwinian population, the dynamics that govern its evolution must respect the symmetries of the RME.

This criterion sanctions some intuitively plausible examples. Consider the group-selection model of GSK introduced in Section 2 when $\alpha_1 = \beta_1$. This would be the case if, for instance, the particle types A and B are two alleles occurring at a common locus and the group types tracked by f_0, f_1, and f_2 are the frequencies of the three possible genotypes. The fact that $\alpha_1 = \beta_1$ in this interpretation reflects the fact that the alleles in a genome (a group) share a common fate. With these assumptions, it is straightforward to show that Equations (2.7) – (2.9) are instances of the diploid RME for three distinct types ($i \in \{0, 1, 2\}$). It is thus necessarily the case that Equations

18

(2.7) – (2.9) jointly respect the symmetries of the RME. So if we take the symmetry approach seriously, the population of groups in this model is in fact a Darwinian population. That is, the genotypes in a diploid population evolve by natural selection given both that the allele frequencies do so and that alleles share a common fate.

It is worth noting that the Symmetry Condition captures cases of both haploid and diploid evolution without any special pleading. This is in contrast with Godfrey-Smith's appeal to 'paradigm populations' (Godfrey-Smith, 2009). Employing a spatial analogy Godfrey-Smith positions each population in a three-dimensional space defined by three parameters: fidelity of heredity, "dependence of realized fitness differences on intrinsic properties," and smoothness of the fitness landscape. Paradigm populations are located in the vicinity of the point maximizing all of these parameters. However, only haploid populations can maximize all three parameters. Thus, we must either treat sexually reproducing populations as somewhat less Darwinian or we have to stretch the region of 'paradigm populations' to include both (Godfrey-Smith opts for the latter). But under the symmetry approach, both types of population can be representations of the fundamental symmetry group that characterizes Darwinian evolution.

Of course, for the Symmetry Condition to be compelling it must rule out some cases of putative evolution by natural selection. Here again we don't have to look too far for an example. To demonstrate that a population governed by a particular dynamics fails to meet the Symmetry Condition, it is sufficient to show that it violates at least one part of the symmetry structure of the RME. Thus, rather than attempt a full explication of the symmetry structure of the RME, it will suffice for our purposes to consider only the most obvious subset of the symmetries of Equation

19

(4.2): fitness scaling. More explicitly, Equation (4.2) is invariant under transformations that change all fitnesses by a multiplicative constant. Thus, a necessary condition for a population to be a Darwinian population is that its dynamics must be invariant under all such fitness-scaling transformations.

Consider again the dynamics of group evolution expressed by Equations (2.7) – (2.9) and this time suppose that $\alpha_1 \neq \beta_1$. There is certainly no reason from the particle perspective why this couldn't be the case. However, if it is the case, then it is straightforward to see that fitness scaling is no longer a symmetry of Equation (8). That is, if the alleles in a group do not share a common fate, then the population of groups does not evolve as a Darwinian population.

6. Conclusions and caveats

I began this essay with a question about the levels of selection: if a population of biological entities evolves by natural selection, is it possible for a population of groups of those entities to do so as well? I suggested that the notion of evolution by natural selection is a dynamic one, and thus we should view the question as one about the sameness of dynamics. This in turn was given a precise meaning in terms of symmetry structures and their representations. In short, two populations whose states form a representation of the same symmetry structure under the relevant set of transformations are said to instantiate equivalent dynamics. Since it is in fact possible for populations of groups to represent the same symmetry structure as the population of entities comprising the groups, we found the answer to our original question to be affirmative.

To be a Darwinian population is to possess a certain sort of dynamical symmetry, and it is possible for populations of groups to possess that sort of symmetry.

There are two important caveats. The first concerns what amounts to a conventional choice, namely a decision about what we should take the fundamental dynamics of evolution by natural selection to be. More to the point, we have to decide what symmetry structure to take as minimal. Perhaps the best approach is simply to view the group of fitness-scaling transformations as the entire symmetry structure that is minimally sufficient for a dynamics to count as Darwinian. The second caveat is that for the symmetry approach to say anything useful about real populations and not idealized infinite ones, it will be necessary to restate the definition of a representation as well as the Symmetry Condition in terms of approximate symmetries. But that is a topic for another time.

Works Cited

Godfrey-Smith, Peter. *Darwinian Populations and Natural Selection*. New York: Oxford
University Press, 2009.

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

Okasha, Samir. *Evolution and the Levels of Selection*. Oxford; Oxford; New York: Clarendon
Press ; Oxford University Press, 2006.

Page, Karen M., and Martin A. Nowak. "Unifying Evolutionary Dynamics." *Journal of
Theoretical Biology* 219 (2002): 93-98.

Rosen, Joe. *Symmetry in Science : An Introduction to the General Theory*. New York: Springer-
Verlag, 1995.

Sober, Elliott, and David Sloan Wilson. "A Critical Review of Philosophical Work on the Units
of Selection Problem." *Philosophy of Science* 61, no. 4 (1994): 534-55.

Sterelny, Kim, and Philip Kitcher. "The Return of the Gene." *The Journal of Philosophy* 85, no.
7 (1988): 339-61.

Wagner, Günter P., and Manfred D. Laubichler. "Character Identification in Evolutionary
Biology: The Role of the Organism." *Theory in Biosciences* 119 (2000): 20-40.

Wilson, David Sloan, and Elliott Sober. "Reviving the Superorganism." *Journal of theoretical
biology* 136, no. 3 (1989): 337-56.