

Kinds of process and the levels of selection

Received: date / Accepted: date

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 of groups. Those approaches to group selection that do emphasize dynamics lack an account of the relevant notion of equivalent dynamics. I show that the theory of ‘dynamical kinds’ I proposed in (Jantzen, 2014) can be used as a framework for assessing dynamical equivalence. That theory is based upon the notion of a dynamical symmetry, a transformation of a system that commutes with its evolution through time. In the proposed framework, structured sets of dynamical symmetries are used to pick out equivalence classes of systems. These classes are large enough to encompass the range of phenomena we associate with natural selection, yet restrictive enough to guarantee a sort of causal homogeneity. By characterizing dynamical kinds via symmetry structures in this way, the question of levels of selection becomes a precise question about which populations respect the dynamical symmetries of Darwinian evolution. Standard population genetic models suggest that populations undergoing evolution by natural selection are partially characterized by a group of fitness-scaling symmetries. I demonstrate conditions under which these symmetries may be satisfied by populations of individuals, populations of groups of individuals, or both simultaneously.

Keywords Darwin, dynamics, evolution, group selection, levels of selection, natural kinds, natural selection, process, symmetry, symmetry transformation, units of selection

1 Framing the problem

This essay concerns the relationship between two biological facts. First, biological systems are hierarchically structured: macromolecules are parts of cells, which are parts of organisms, which are parts of demes, and so on. Second, at least some biological systems undergo evolution by natural selection. To ask how the occurrence of natural selection relates to the structural hierarchy is to ask about the so-called ‘levels of selection’. Much has been written on this subject, and over decades of discussion, a tangled bank of problems has grown up around this theme.¹ These include questions such as: On which biological entities in the structural hierarchy does natural selection act (Vrba and Gould, 1986)? Which entities possess a fitness (Brandon, 1982)? At what level of the hierarchy do we find individuals (Wilson and Sober, 1989)? Which entities replicate (Dawkins, 1976)? What entities benefit from selection (Sober and Wilson, 1994)? At what level does adaptation occur (Sober, 1993)? Can selection take place at multiple levels simultaneously (Damuth and Heisler, 1988)?

Despite the apparent diversity, all of these questions are to some extent derivative of a more basic concern: at what levels in the biological hierarchy is the process of natural selection to be found? The reference to “process” here is not coincidental.² Rather, the notion of process is conceptually central. Whatever else it

¹ See Okasha (2006b, 2011) for an overview of the historical arc of the debate, and Lloyd (2012) for a survey of some of the conceptual issues raised.

² Nor is it unorthodox. In his consideration of the evidence for or against selection at various levels of the structural hierarchy, Hull (1980, pg. 312-3) proposes “...to investigate the general characteristics of the evolutionary *process* at some length and then

may be, evolution by natural selection is a process. In fact, it is an instance of a special sort of process which, for lack of a better term, I'll call a *dynamical process*. Processes in general involve change through time of the state of a system. In a dynamical process, the state of the system corresponds to the values of a collection of causally connected variables. In such a system, not all states are possible, and which are accessible is determined by causal relations amongst the variables. To illustrate the contrast with processes in general, consider the change through time of the system consisting of all the coins in my left pocket along with those in the left pocket of someone in Beijing. Any number of coins in my pocket is compatible with any number of coins in my Chinese counterpart's pocket. The 'system' in question does evolve through time and so instantiates a process. But it is not a dynamical process. In contrast, the changing state of a chemical reaction described in terms of the concentrations of reagents and products is a dynamical process—changing the value of one variable changes the values that are accessible to all variables, at that time or in the future.

Evolution by natural selection is a dynamical process in this sense, one that involves the change through time of type frequencies in a population. The central question of the levels of selection is a question about the conditions under which a population of biological entities can be said to instantiate a dynamical process of the Darwinian kind. This in turn demands a theory of kinds of dynamical process, or *dynamical kinds*—an account of how particular processes sort into natural and distinct kinds. Only with such a theory can it be determined whether a process at one or another level of the structural hierarchy is an instance of evolution by natural selection or of a different sort of process altogether.

Standard approaches to the levels of selection tend not to focus on identifying characteristics of the process of natural selection but rather on identifying static conditions necessary and sufficient for it to occur. Arguably, this is what Darwin was up to in the first four chapters of the *Origin of Species* (1993), and it is certainly what Lewontin (1970) had in mind when he proposed his tripartite schema.³ According to the latter, a population of entities (biological or otherwise) will undergo evolution by natural selection if: (i) there is phenotypic variation amongst members of the population; (ii) there is differential fitness among the phenotypes; and (iii) there is correlation in fitness between parents and offspring (i.e., fitness is heritable). To determine whether or not these conditions are or can be met by populations of entities at various levels of the structural hierarchy, one must first clarify the notions of heritability and reproduction. Unsurprisingly, those who take Lewontin's approach have developed these notions in great detail. This is, for example, the bulk of the work taken up by Godfrey-Smith (2009) in his exemplary monograph on 'Darwinian populations'.

From the process perspective, however, Lewontin's properties-first approach is question-begging. Evolution by natural selection is a dynamical process of populations. The process itself—irrespective of whatever mechanisms underly it—can be described solely in terms of variables that refer to population properties and how they change through time. Put differently, it is possible to characterize the process without assuming anything about the nature and characteristics of the participating populations other than that they possess properties of type frequency and type fitness. It is not necessary to know what properties of biological entities are in fact used to instantiate such a process. Assuming at the outset facts about those properties—for instance, that members of an evolving population must reproduce in a particular fashion thus begs the question about which populations can instantiate a process of the Darwinian kind. To put a sharper point on the criticism, we cannot know whether the conditions Lewontin offers are in fact necessary and sufficient unless we already have a way to decide which processes are positive instances of natural selection.

For these reasons, I suggest that we consider what the levels of selection problem looks like if we set aside the properties-first approach of Lewontin and his intellectual descendants, and instead focus on directly characterizing the dynamical process of evolution.⁴ For this, we need a theory of dynamical kinds, and I will tentatively adopt one of my own (Jantzen, 2014). As we'll see, the result is a framework for asking and precisely answering questions about the interaction of kinds of processes with the structural hierarchy. This framework is consonant with recent work stressing the causal distinctions between individual selection (specifically kin selection) and group selection (Okasha, 2016; Krupp, 2016). But it also makes some novel distinctions, and supports a stronger pluralism than is sanctioned by those who see levels of selection as mere descriptive choices. Of course, my account is not the only conceivable way to spell out a theory of dynamical

to discuss only briefly the particular entities that may or may not possess the characteristics necessary to function in this *process*" (emphasis added). More recently, Godfrey-Smith (2009, p. 109) dedicates a chapter to "...the idea that Darwinian *processes* occur at different "levels" (emphasis added).

³ See also (Godfrey-Smith, 2007).

⁴ I intend this not as a demand, but rather a plea to try out a heretofore unexplored perspective.

kinds, and it may turn out that there are compelling reasons for adopting a different view. In that case, this essay serves as a template for developing process-first views. A key component of that template is a theory of kinds of process or ‘dynamical kinds’.

2 An incomplete framework

What must a theory of dynamical kinds look like if it is to serve its intended purpose with respect to the levels of selection question? To answer this, it will be instructive to examine a prominent framework for discussing multilevel selection that explicitly emphasizes dynamics. I have in mind the formalism developed by Kerr and Godfrey-Smith (2002) to talk about evolutionary processes at both the individual and group level. Despite the merits of their framework, it does not provide the resources to answer the question of levels of selection as construed above. Understanding why this is the case will suggest general conditions on the sort of theory of dynamical kinds that is needed.

2.1 The KGS model

Kerr and Godfrey-Smith (hereafter KGS) ask us to consider an infinite population of biological entities that are of two types, A and B. These entities, whatever they may be, I’ll call ‘particles’ as KGS do in order to avoid confusing and prejudicial use of terms like ‘individual’ or ‘organism’. These particles are supposed to aggregate into groups, all of which contain n particles. Each group is classified according to 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 – particles replicate within groups through simple copying, groups dissolve into a single, panmictic mass of particles, the old particles die off, and then new groups form from the particle offspring.

According to KGS, 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. Specifically, let α_i be the fitness of an A particle in a group with i A’s and β_i stand for the fitness of a B particle in a group with i A’s. In this approach then, groups are treated as local environments across which particle fitnesses vary.

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

$$p(t+1) = \frac{1}{\bar{w}} \sum_{i=1}^n f_i(t) \frac{i}{n} \alpha_i \quad (1)$$

$$q(t+1) = \frac{1}{\bar{w}} \sum_{i=0}^{n-1} f_i(t) \frac{n-i}{n} \beta_i \quad (2)$$

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

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

Equations (1) and (2) are not dynamically sufficient—they do not determine the time-course of the frequencies p and q —unless we specify the functions $f_i(t)$. In the cases considered by KGS, the $f_i(t)$ are assumed to be functions solely of $p(t)$.

In constructing (1) and (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. KGS claim that we can do so in the following manner. To each group we assign two parameters:

$$\begin{aligned}\pi_i &= \text{total number of particle copies from a group with } i \text{ A types} \\ \phi_i &= \frac{(\text{number of A copies in a group with } i \text{ A types})}{(\text{total number of copies in a group with } i \text{ A types})}\end{aligned}$$

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:

$$\bar{\pi}p(t+1) = \sum_{i=1}^n f_i(t)\phi_i\pi_i \quad (3)$$

$$\bar{\pi}q(t+1) = \sum_{i=0}^{n-1} f_i(t)(1-\phi_i)\pi_i \quad (4)$$

In both of these equations,

$$\bar{\pi} = \sum_{i=0}^n f_i(t)\pi_i$$

2.2 Group dynamics in KGS

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 (3) and (4) represent some facts about groups, a glance at the definitions of π_i and ϕ_i reminds us 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 (3) and (4) 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 (1) and (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:

$$f_i(t) = \binom{n}{i} (p(t))^i (1-p(t))^{n-i} \quad (5)$$

For simplicity, I will assume that $n = 2$. That is, each group is composed of just two particles. From Equations (1) – (5) we can then find explicit expressions for the dynamics of the group frequencies:

$$f_0(t+1) = \frac{(\beta_0 f_0(t) + \frac{1}{2}\beta_1 f_1(t))^2}{(\alpha_2 f_2(t) + \beta_0 f_0(t) + \frac{1}{2}(\alpha_1 + \beta_1) f_1(t))^2} \quad (6)$$

$$f_1(t+1) = \frac{2(\frac{1}{2}\alpha_1 f_1(t) + \alpha_2 f_2(t))(\beta_0 f_0(t) + \frac{1}{2}\beta_1 f_1(t))}{(\alpha_2 f_2(t) + \beta_0 f_0(t) + \frac{1}{2}(\alpha_1 + \beta_1) f_1(t))^2} \quad (7)$$

$$f_2(t+1) = \frac{(\frac{1}{2}\alpha_1 f_1(t) + \alpha_2 f_2(t))^2}{(\alpha_2 f_2(t) + \beta_0 f_0(t) + \frac{1}{2}(\alpha_1 + \beta_1) f_1(t))^2} \quad (8)$$

To reiterate, what Equations (6) – (8) represent are the dynamics of a population of groups as it changes in time, assuming that the particles composing the groups are governed by Equations (1) and (2).⁵ These expressions, not Equations (3) and (4), are the group 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 4.2 and Appendix A below) to identify group fitnesses in this case with the coefficients of the frequencies appearing in the denominators of each expression on the right-hand side, namely α_2 , β_0 , and $\frac{\alpha_1 + \beta_1}{2}$. These are proportional to the parameters π_i introduced as ‘group fitnesses’ by KGS. But the identification of group fitnesses cannot be determined by simple inspection of the equations. Only the group frequencies are unambiguously recognizable. What’s more, the form of, say Equation (6) is quite different from that of Equation (1) or (2). In particular, the frequency of a particle type at time $t + 1$ depends on only one particle frequency at t whereas each group frequency at time $t + 1$ is a function of two frequencies. The dynamics governing change in the population of groups looks quite different from the dynamics of particles.

2.3 Isn’t it just a matter of perspective?

A number of authors have defended the claim that there is no fact of the matter whether selection occurs at the level of the group or of the individual particle – the distinction is merely one of descriptive perspective (Kerr and Godfrey-Smith, 2002; Sterelny, 1996; Dugatkin and Reeve, 1994). According to this ‘descriptivist’ view, we are free to adopt the particle perspective or the group perspective as we wish to describe what amounts to one and the same underlying evolutionary process. Consequently, we should view Equations (1)–(2) and Equations (6)–(8) as saying the same thing.⁶

There are at least two reasons that weigh against the descriptivist view. The first is asymmetry of information.⁷ One might suppose that a necessary condition for the equivalence of descriptions is inter-translatability. In other words, if we have available two frameworks that allow us to construct ‘equivalent’ descriptions of a phenomenon, then we ought to be able to infer one description from the other and vice versa. In the example above, we were able to infer a group-level description from the particle-level description, but only because: (i) group types were defined by their particle composition, (ii) we assumed that particle fitnesses are exactly determined by group membership, and (iii) we assumed a particular relation between the f_i ’s and p (see Equation 5). Suppose we keep the first two assumptions – that group types are defined in terms of particle composition and particle fitnesses are determined by group membership – but make no assumption about the connection between particle and group frequencies. Then in general, if we were given only a dynamically sufficient particle-level description of an evolving population, we could *not* infer the group-level description. This follows from two facts. First, the constraints imposed by the definition of group types are generally insufficient to uniquely determine a distribution of group types given the distribution of particle types. And second, if the particle fitnesses are not sensitive to all group distinctions (e.g., if $\beta_0 = \alpha_2$ in the KGS model), then only a non-invertible function of group frequencies needs to be known to make the particle model dynamically sufficient. Thus, knowing the particle model does not entail the group model. Similarly, there is no a priori reason to insist that the f_i be functions of p . If they are not, then we cannot infer the particle-level description from the group-level description either. At least, we cannot infer a predictive particle-level dynamics from the group description, though we could, of course, track the particle frequencies using the group type frequencies.⁸

Now drop our first two assumptions as well, and suppose that group types are characterized in a manner that is logically independent of particle frequencies and that particle fitnesses are not exactly determined by group membership. This is, in many circumstances, biologically realistic. We may, for example, be interested in counting ant colonies and classifying them into group types in terms of, say, the ratios of different caste subpopulations. Those group type definitions do not entail anything about allele (particle) frequencies, either

⁵ Equations (3) and (4) represent what Damuth and Heisler (1988) refer to as multilevel selection [1] (MLS1 in (Okasha, 2006a)), while Equations (6) – (8) represent the multilevel selection [2] perspective (MLS2 in (Okasha, 2006a)).

⁶ Wilson (2003) calls this position “model pluralism” (and argues against it).

⁷ A related and more extensive argument for informational asymmetry can be found in Lloyd et al. (2008).

⁸ Even this much may be impossible if not all particles join groups.

within colonies or across the aggregate of colonies. In such a case, all bets are off. In scenarios like this, it is a contingent empirical fact how the group and particle frequencies are related, such that knowing one is insufficient to determine the other. Since it is at least possible that real populations instantiate dynamics with one or another of these informational asymmetries, it matters which perspective one adopts, at least insofar as one wishes to accurately model the evolutionary process.

The second reason for doubting the equivalence thesis is causal asymmetry. This is reflected in the informational asymmetry described above. Consider those cases in which group frequencies do not uniquely determine particle frequencies, and vice versa.⁹ It may be the case, for example, that the future distribution of group frequencies depends causally only upon the current group frequencies and not directly upon the particle frequencies. Thus, keeping the particle frequencies fixed but changing the group frequencies will change how both the group and particle populations change through time. In this case, it looks like selection is really occurring at the group level. On the other hand, it may be that interventions on group frequencies that leave particle frequencies intact are totally irrelevant for determining future states. In that case, it's only the particle frequencies that are causally salient. This is exactly the sort of causal asymmetry that Okasha (2016) has pointed out in arguing against the descriptivist thesis with respect to kin selection (a particular variety of structured particle selection) versus group selection. In principle then, it's not merely a matter of description. It may be that no real populations ever instantiate an evolutionary process for which the descriptions aren't equivalent, but since we can't know that in advance, we must accept the possibility that only one or the other perspective is correct with respect to accurately reflecting the causal situation.

That said, even if both perspectives are apt in a given case – even if descriptions at the group and particle level are equivalent in the informational sense and both capture something of the causal structure of the world – it may be that the process described at one level is dramatically different in kind from that occurring at the other level. That is, we have no reason to suppose that just because natural selection occurs at the level of particles and the group level description is causally apt, that natural selection is therefore taking place at the group level. There are many ways in which processes at different levels can come apart. This seems to be a general feature of supervenient processes, if indeed there be any. For example, the individual growth of a tree in terms of the increasing biomass of the high-level organism looks quite different from the process of cellular population growth that underlies it. The former exhibits allometries not obvious in models of the latter, and the latter involves competition amongst subunits and cell lines that is not part of the tree-level description. While both are causally and descriptively apt, the models one would use at the tree versus cell level are qualitatively quite different, with different scope and supporting distinct varieties of inference. The cell model, for instance, might also be applied to independent cells in a biofilm where the high-level model may not have any suitable parameterization (what counts as the 'tree'?). In short, choice of descriptive level matters, as does the decision to call the dynamics at either level evolutionary.

2.4 What's missing

To take stock so far, the KGS formalism supports the construction of models of particle evolution for which groups manifest as fitness structure. With the additional assumption of a particular connection between particles and groups, we were also 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. The result was a set of dynamical relations that look superficially quite different from those governing the particles. In this context the levels of selection question amounts to this: which, if either, dynamics is an instance of natural selection? Is the dynamics expressed by either model of the right sort to count as evolution by natural selection? There is no answer inherent in the KGS framework. Nor is there any a priori reason to think the answer is the same for particles and groups. Even if we assume that the particles are undergoing evolution by natural selection, we have no principled means of deciding whether the dynamics expressed in Equations (6)–(8) is “the same as” the dynamics expressed in Equations (1) and (2) in some relevant sense. Of course, the hard part is specifying the “relevant sense”. What is it to be the same kind of

⁹ The reason for this restriction is that asserting causal relations amongst variables is problematic when some of those variables are related by logical entailment. For illuminating discussions of the problem, see (Spirtes, 2009; Woodward, 2015). But if we stick to the scenarios in which group types and frequencies are independent of particle frequencies in the sense that it is possible to intervene on one while holding the other fixed, no such problems arise.

dynamical process? Note that I am not asking whether either the population of particles or groups 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. Even if there is such a mechanism, it is insufficient to guarantee any particular dynamics, and it is obeying a dynamics of the right sort that makes a dynamical process one of natural selection.

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) reformulate the question, “What is an evolutionary trait?” as a question about dynamics: “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, like KGS, Wagner and Laubichler offer no account of “sameness of dynamics”, though they appeal to the concept. 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 (1) and (2) above. But even the addition of a third allele or just one more fitness environment would be described by a different set of equations, as would frequency-dependent fitnesses. Obviously, the ‘sameness’ of dynamical equations should not be taken too literally. But how can we decide whether or not the dynamics of two changing populations is the ‘same’ in some more liberal, salient sense? And how can we decide if either is evolution by natural selection?

An answer to the question of sameness can only be provided by a theory of dynamical kinds. Specifically, what we need is a theory of dynamical kinds that is permissive enough to treat systems of differing algebraic form (e.g., systems with and without frequency dependent selection) as belonging to the same kind, that can equate dynamics involving differing numbers of dynamical variables (e.g., evolving populations with two versus three types), and yet nonetheless ensures that ‘equivalent’ dynamical systems constitute something at least approximating a natural kind. In the next section, I present an account with all of these features: my theory of dynamical kinds (Jantzen, 2014).

3 A theory of dynamical kinds

3.1 Motivating the approach

My account is built upon the notion of a ‘dynamical symmetry’. This is not an obvious choice of feature with which to individuate kinds of process. So before I present the technical details, it may be helpful to motivate the approach by considering some simple biological examples.¹⁰ Since I do not want to beg the question with respect to levels of selection, I’ll draw upon a different process of populations: how they grow. How does the number of individuals in a local population change with time? There is, of course, no unique answer. Population size, for instance, follows some approximation of an exponential curve in time given appropriate conditions. In other conditions – e.g., limited space and resources – populations exhibit rapid growth followed by what looks like an asymptotic approach to a finite “carrying capacity”. In still other circumstances, populations are seen to oscillate or even vary chaotically in time.¹¹ And for any one of these general patterns, there are myriad variations. Despite all of this diversity, can we say anything about *kinds* of growth?

To make the example more concrete, suppose we are concerned with two local, possibly interacting populations of organisms belonging to different trophic species. For such a pairing, all of the behaviors mentioned above are possible. What can be said about kinds of behavior for this two-species system over time, or in other words, about kinds of two-species growth process? Perhaps the most obvious way we might proceed is to stipulate a regression model. That is, we could stipulate a pair of parametrized functions for describing the populations over time (one for each species in the system). Then any other two-species system would be an instance of the same kind of growth process just if we can find a set of parameter values for our

¹⁰ This is not how I motivate the account in (Jantzen, 2014). Rather I argue for accepting dynamical kinds as a (partial) solution to what I call the epistemic question of natural kinds: what, if anything, do categories that sustain inductive generalization have in common? My examples in that paper are drawn from a wider range of scientific disciplines, including physics and chemistry.

¹¹ For an overview, see, e.g., (May and McLean, 2007). In all realistic circumstances, growth is stochastic.

model such that the model describes the observed growth in both systems. In other words, if our regression model can represent both two-species growth processes given a judicious selection of parameter values, then the growth processes are instances of the same kind.

But this strategy would seem to be both too restrictive and too permissive. It's restrictive in that it cannot accommodate the sorts of variety we've been talking about. For example, physicists certainly talk as if the damped harmonic oscillator is a dynamical kind. At least, it's plausible that all masses oscillating on springs with varying degrees of frictional damping are instances of the same kind of process. After all, they are all described (albeit with differing parameter values) by differential equations of the form: $\ddot{x} + 2\zeta\omega_0\dot{x} + \omega_0^2x = 0$. But for some such systems, the mass will oscillate periodically, for others, it will drop to its minimum position in what looks like exponential decay. In other words, had one chosen a parameterized sine wave or exponential to characterize the kind, one would have arbitrarily excluded dynamical processes that share a great deal with those in the kind. Similarly, for any reasonably complex model of two-species growth – such as that expressed by the competitive Lotka-Volterra equations (see, e.g., (Pastor, 2011)) – a great variety of specific trajectories through time are possible. It would be hard to choose a single pair of parameterized functions, and any such choice would of necessity be arbitrary – it would separate processes that in fact share deep similarities. We would judge oscillating populations of pairs of fish trophic species¹² as different from stable pairs of competing plankton populations despite the fact that both systems are described by, say, identical differential equations modulo a few parameter values.




Fig. 1 (a) and (c): the populations of system A over time for two different initial conditions; (b) and (d): the populations of system B over time for two different initial conditions.

On the other hand, the approach of stipulating a particular parametric model of change through time is too permissive in that causally disparate processes would be lumped together. Even if two systems of populations belong to the same kind in this sense, they may behave quite differently with respect to perturbations of their state. For example, consider the extreme case depicted in Figure 3.1 (a) in which two populations (let's assume they belong to two distinct trophic species) vary in the simplest possible way: they stay fixed through time. Call this system A. There are myriad parameterized regression models we could identify that, for a suitable choice of parameters, emulate the behavior of system A perfectly well. Let's call the model we've chosen M . Since the curves in Figure 3.1 (b) showing the behavior of system B are identical to those in Figure 3.1 (a), we must concede that both system A and system B belong to the same dynamical kind – they are both instance of the 'same' process because there are parameterizations of M that capture both. In fact, in this case, the same parameter values work for both. But A and B differ causally. Figure 3.1 (c) shows what happens to system A when we halve the initial populations. Not much happens at all – both populations remain constant through time. But in Figure 3.1 (d) you can see that reducing the initial populations of system B by a factor of 2 has a dramatic qualitative effect – the two species in this system

¹² This really does happen: see (Hamrin and Persson, 1986).

do in fact interact with one another, and their populations now vary through time.¹³ The point is that the causal connections that generate a process can differ from system to system even though it is possible to find instances of each sort of causal system that have observationally indistinguishable trajectories. If we stipulate a regression model to define dynamical kinds, then some of the time we'll sweep together processes that are only superficially similar.

How could we avoid such an obvious deficiency? We might try to generalize the approach by stipulating a whole *class* of parameterized functions describing changes through time rather than a particular function. For example, we might identify sameness of dynamics with “being a solution of the same differential equation”. That way, all of the damped harmonic oscillators would properly belong to a single kind of process, as would all the ways by which populations governed by the Lotka-Volterra equations can evolve through time. But this wouldn't help us much in distinguishing A from B in the example above. That's because the zero-growth, non-interacting model can be obtained by setting certain parameters to 0 in the competitive Lotka-Volterra equations which in turn describe the non-trivial growth of System B.¹⁴ We would have to invoke some rule for precluding these parameter values while allowing others in. In effect, this approach just forces us to confront anew our original question: what counts as the same differential equation? Do different parameter values matter? In other words, is $\dot{x} = x$ the same as $\dot{x} = 2x$? Or $\dot{x} = x^2$? What about dimensionality? Does the competitive Lotka-Volterra system of equations in two variables (for two species) describe the same kind of process as the Lotka-Volterra system in three variables? Or 25? After all, qualitatively different phenomena can arise with increasing dimension. This route doesn't look very promising for getting a handle on growth, let alone providing a general theory of kinds of process.

But there's another way to go: there are some generic properties of causal structure that pick out rich classes that are neither overly restrictive nor overly promiscuous. Let's further simplify our example and consider only the growth of a single population (a population belonging to a single trophic species). Suppose that the categories of 'logistic growth' (bounded growth for which $\dot{x} = rx(1 - x/K)$) and 'exponential growth' (according to $\dot{x} = rx$) are at least approximately the sort of categories we're looking for. Then one set of unambiguous and empirically accessible features of the dynamics that distinguishes these kinds of growing population involve a sort of indifference with respect to intervention. Suppose a population is growing exponentially. There are lots of distinct growth curves that answer to such a description, one for every possible combination of intrinsic growth rate, r , and initial population size, x_0 . But despite this variety, there is a simple universal feature of all exponentially growing populations, and it is this: if you scale the initial population from x_0 to kx_0 and then let it grow over some time interval Δ , the resulting population will be exactly what you'd have seen had you let it grow from x_0 and then, after waiting through the interval Δ , multiplied the population by k . It doesn't matter what Δ one chooses, or what x_0 is, or what the intrinsic growth rate, r is; this is a universal feature of populations that grow according to a first order differential equation of the form $\dot{x} = rx$. Importantly, this indifference to scaling is *not* a feature of populations whose growth is governed by, e.g., the Verhulst logistic equation, $\dot{x} = rx(1 - x/K)$. Quite generally for a population exhibiting logistic growth, it matters whether you first scale the population and then allow it to grow unmolested, or grow first and then scale – you get different results depending upon the order.

3.2 Symmetries

So when can two growing populations be said to exhibit the same kind of growth process? The suggestion to be drawn from the preceding discussion is that in order to answer this question 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, σ , of the system that is invisible to the dynamics that govern that system in the sense that it makes no difference whether we apply σ and then advance the state through time according to the dynamical laws of that system, or if we instead apply the laws to advance the state and

¹³ The model used to generate data for system A is a zero-growth model: $\dot{x}_i = 0$. For system B, it was assumed that the populations obey the two-species competitive Lotka-Volterra model in which $\dot{x}_1 = r_1x_1(1 - (x_1 + \alpha_{1,2}x_2)/K1)$ and $\dot{x}_2 = r_2x_2(1 - (x_2 + \alpha_{2,1}x_1)/K2)$.

¹⁴ Specifically, $r = \alpha_{i,j} = 0$.

then transform the result with σ . Either way we end up with the system in the same state.¹⁵ This idea can be put more precisely as follows (Jantzen, 2014, p3632):

Definition. Let S be the set of states of a system and let $A_{t_0, t_1} : S \rightarrow S$ be the time-evolution operator which takes the system from state S_0 at t_0 to S_1 at t_1 . A *dynamical symmetry* is any operator $\sigma : S \rightarrow S$ with the following property: $\forall s \in S \forall t_0 \forall t_1 > t_0 [A_{t_0, t_1}(\sigma(s)) = \sigma(A_{t_0, t_1}(s))]$.

As I indicated above, scaling by a constant (any positive constant) meets this condition for populations in exponential growth, or in other words, populations whose growth is governed by the simple differential equation, $\dot{x} = rx$ (where r is a positive constant representing intrinsic growth rate). Though multiplicative scaling is *not* a dynamical symmetry of logistic growth (growth according to $\dot{x} = rx(1 - x/K)$), there is a set of transformations that meet the requisite condition. Specifically, if K is the carrying capacity for such a population, then for every real value of p , the following transformations commute with growth through time:

$$\sigma_p(x) = Kx / ((1 - e^{-p})x + e^{-p}K)$$

You end up with the same population size whether you first transform the population according to σ_p (for some particular value of p) and then let it grow, or let it grow and then apply σ_p .

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 $\sigma_2 \circ \sigma_1$, 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.¹⁶ A particular group structure along with a list of the specific, concrete transformations that manifest that structure constitute the *symmetry structure* of a dynamics:

Definition. A *symmetry structure* is a collection of physical symmetry transformations, represented by $\Sigma = \{\sigma_i | i = 1, 2, \dots\}$, along with a composition function $\circ : \Sigma \times \Sigma \rightarrow \Sigma$ that indicates how those symmetries behave under composition. (For a discrete set of symmetries, the composition function can be given in the form of a multiplication table.)

It is important to emphasize that symmetry structures are *not* abstract algebraic objects. Rather, a symmetry structure is a collection of concrete, physical transformations whose behavior under composition is described by a particular abstract algebraic object, i.e., a group. To make this point as explicit as possible, consider the close analogy between electric circuits and hydrostatics. So close is the analogy in behavior that circuits have been used to model complex hydrostatic systems with great accuracy (see Bourouina and Grandchamp, 1996). This is because the dynamics of certain sorts of electric circuits and the dynamics of certain hydrostatic systems are described by the same equations; the mathematical variables are simply interpreted differently. Though they are described by identical equations, however, these two systems have different symmetry structures in the sense in which I use the term. This is because the transformations for hydrostatic systems involve operations like changes in pressure while the transformations for circuits involve completely distinct operations like changing voltage. Their symmetry structures share a common abstract group structure, but they are not identical because the concrete, physical transformations for each are different.

3.3 Dynamical kinds

Symmetry structures offer a way of classifying individual processes on the basis of their dynamical properties (Jantzen, 2014, p3635):¹⁷

Definition 1 (Dynamical kind) A *dynamical kind* is a class of systems of variables that share a set of dynamical symmetries that are related by a non-trivial symmetry structure.

¹⁵ This is what I call a ‘dynamical symmetry with respect to time’ in (Jantzen, 2014). The concept of a dynamical symmetry simpliciter is rather broader than this, and applies to atemporal causal systems.

¹⁶ Elsewhere (Jantzen, 2017b, p20), I have suggested that 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 monoids.

¹⁷ By ‘non-trivial’, I mean that it is not the case the all transformations are dynamical symmetries as would be so if none of the variables in the system are causally related to one another.




Fig. 2 The solid line shows an exponential curve, a solution to $\dot{x} = rx$. The dashed line shows a solution to the equation $\ddot{x} + \dot{x} = rx$. Systems governed by either equation belong to the same dynamical kind.

For a particular system to instantiate a particular kind of dynamics then, it is necessary and sufficient that the system manifests the symmetry structure associated with that kind. So, for instance, if a system is to be a Newtonian gravitating system, then it must have the dynamical symmetry structure of Newton’s laws of motion and gravitation. In particular, the dynamics must be insensitive to rigid translations and rotations of the system (and a handful of other such transformations) as well as composites of these transformations.

Prima facie, this might all 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 some of the systems we want to lump together as Darwinian evolvers do not obey exactly the same dynamical 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 generally a sufficient condition for that system to exhibit a particular dynamics – two systems that respect the same symmetry structure might differ in some dynamical details.

To return to biological growth, consider a single population with growth described by $\ddot{x} + \dot{x} = rx$. The possible growth curves for this population are quite distinct from those of exponential growth (see Figure 3.3; the dashed line traces a solution to this differential equation). Yet both ordinary exponential growth and this new sort of growth belong to the same dynamical kind, namely that whose symmetries are of the form $\sigma_k(x) = kx$ and for which $\sigma_{k_1}(x) \circ \sigma_{k_2}(x) = \sigma_{(k_1 \cdot k_2)}(x)$. Put qualitatively, this is the class of scale-invariant growth, or growth for which absolute population size is irrelevant. It contrasts with growth characterized by, say, a fixed carrying capacity. Scale invariant growth may or may not be a biologically relevant kind – I don’t know if anything actually grows in a scale-invariant way. But it’s clear that all such systems would share some significant features in common, features that make recognizing the kind worthwhile. But they would not share *all* of their features.

What about systems involving different numbers of variables? Here my earlier account (Jantzen, 2014) is a bit ambiguous. Symmetry structures are defined with respect to a particular set of variables, so unless two systems share all of the variables constitutive of the symmetry structure (and thus have the same number of variables), it would seem that they cannot belong to the same kind. Thus, a system with only one population variable or allele frequency cannot belong to a dynamical kind involving two populations or two allele frequencies. But that’s an unnecessarily strict view of what constitutes a variable, and contradicts what I’ve said about, e.g., physical systems. This is obviously an extremely broad dynamical kind, but it’s Eugene Wigner’s approach to characterizing physical systems that inspired my approach. At any rate, I have asserted, for instance, that “all rigid translations, rotations, and Lorentz boosts are dynamical symmetries of physical systems, with respect to time” (2014, p3637). I didn’t say “physical systems with three position variables” or “physical systems of exactly 5 gravitating masses” or “rigid translations in exactly two dimensions.” Rather, I just said “physical system”. I implicitly assumed that systems involving differing numbers of variables can nonetheless share a common symmetry structure, one that involves more variables in some systems than in others.

How could this work? If we’re going to have a complete enough theory of dynamical kinds to approach the levels of selection, we need to make explicit what I previously left implicit. Here’s one possibility. If we allow for vector-valued variables – in particular, if a variable can take on vectors of varying dimension as arguments – then we can accommodate differing numbers of variables within a single dynamical kind. Here’s what I mean. Consider two systems I’ll call Single and Double. In Single, there is a population of one trophic species

growing according to $\dot{x} = kx$. In Double, there are two trophic species whose growth is well-represented by a system of differential equations: $\dot{x} = k_1(x + y)$ and $\dot{y} = k_2(x + y)$. Double has a single-parameter family of dynamical symmetries in which $x \rightarrow px$ and $y \rightarrow py$. In other words, scaling both x and y simultaneously is a dynamical symmetry of Double. But such a transformation is not, strictly speaking, available for Single since it has no variable y . Of course, Single also exhibits a one-parameter family of dynamical symmetries that involve scaling by a constant (see the discussion of exponential growth above). But now let's describe the situation in terms of population vectors. That is, both Single and Double may be viewed as involving only a single, vector-valued population variable $\mathbf{x}$. In Single, $\mathbf{x}$ has only one component; in Double it has two. Expressed this way, we can identify a single, one-parameter family of vector functions (functions mapping vectors to vectors) comprising symmetries of *both* Single and Double: $\sigma_p(\mathbf{x}) = p\mathbf{x}$. In other words, treating variables in disparate systems as vector-valued (where the dimension of the corresponding vector space varies from system to system) and specifying interventions on the systems via vector functions allows us to see Single and Double as belonging to the same dynamical kind. Doing so captures a deep and genuine similarity between the two – a sort of scale invariance or indifference to absolute population size. At any rate, this approach is consonant with what I've said before.

3.4 Why dynamical kinds based on symmetry?

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? For our present purposes, the most salient reason is that dynamical kinds as defined above offer an attractive solution to the question of 'sameness of dynamics' with which we began. Systems with different numbers of dynamical variables (e.g., different numbers of allele frequencies) may belong to the same dynamical kind, as can systems described by dynamical equations of differing form (e.g., models with frequency-dependent fitnesses versus models with frequency-independent fitnesses). Each process in a dynamical kind evolves through time with an indifference to exactly the same set of perturbations (the same set of dynamical symmetries). Consequently, dynamical kinds are categorically distinct and objective features of the world. Whether or not they constitute the sole answer to the question of natural kinds, they are natural enough to play the role demanded of them here.

This should be reason enough for us to take the approach seriously, since we only need a defensible theory of sameness-of-process. But in (Jantzen, 2014), I defend the account as a general theory of (at least an important subset of) natural kinds. That is, I argue that dynamical kinds are projectible kinds, and that given this is the case, the theory offers a way to explain how it is we recognize projectible kinds with such high efficiency. In the first place, I attempt to demonstrate that dynamical kinds do a better job of capturing many or most scientifically salient kinds than any competing theory of natural kinds. In the second place, I argue that "in a sense, this account of natural kinds is testable; it suggests adopting particular methods for efficient induction in scientific practice. Insofar as these methodological consequences bear fruit, the dynamical kinds approach is that much more plausible..." (Jantzen, 2014, p3644). It would lead us far astray to try to assess the empirical case for taking dynamical kinds seriously as natural kinds. But it seems enough progress has been made on this front to lend weight to the account (see (Jantzen, 2017a)).

4 Dynamical kinds and the levels of selection

4.1 Darwinian dynamical kinds

So what is a Darwinian process? 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 symmetry structure of evolution by natural selection – in which case selection does occur at that level – or it violates this structure, in which case selection cannot be said to operate at that level. To ascertain whether or not a particular population is a Darwinian evolver – a population evolving through a process of natural selection – we need only ascertain whether the dynamics of the population obeys the symmetry structure of evolution by natural selection. Of course, this means addressing the daunting question

of what the relevant symmetry structure is. What are the symmetries of evolution by natural selection, and how do they compose?

There are two ways we might proceed, one empirical and one theoretical. Empirically, we could examine the sorts of transformations of type frequencies that commute with time evolution in real evolving populations. Of course, we won't know in advance which processes to count as evolutionary. Rather, we would hope to see that populations exhibiting changes of type frequency with respect to time fall into a handful of dynamical kinds, one of which more or less aligns with what we've historically identified as evolution by natural selection. But for a variety of reasons, this strategy is impractical at present. The other approach is to look for one or a handful of theoretical models that are considered well-justified and which capture all or most of the systems uncontroversially thought to be undergoing evolution by natural selection. We could then use the symmetry structure of that theory as a tentative candidate for characterizing evolution by natural selection. Restricting our attention to deterministic models of population dynamics (which generally assume infinite populations), the list of models that fit this description is topped by the so-called 'Replicator-Mutator Equation' or RME. This equation comes in different forms for haploid and diploid populations:¹⁸

The RME for discrete generations (haploid):

$$x_i(t+1) = \frac{1}{\bar{w}} \sum_j q_{ji} w_j x_j(t) \quad (9)$$

The RME for discrete generations (diploid):

$$x_i(t+1) = \frac{1}{\bar{w}} \sum_{jk} x_j(t) x_k(t) R_{jk} Q_{jki} \quad (10)$$

The variable x_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_{ji} is the probability that an individual of type j produces an individual of type i 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 (10), $R_{jk} = R_{kj}$ is the expected number of offspring produced by pairs of individuals of types j and k , and so $w_i = \sum_j x_j R_{ji}$. In both cases, $\bar{w} = \sum_i w_i x_i$.

Most if not all of the standard models in population genetics can be cast as special cases of either Equations (9) or (10). In fact, the haploid RME can be treated as a special case of the diploid RME for which $\sum_k x_k R_{jk} Q_{jki} = w_j q_{ji}$. The Price Equation, which describes the change in mean character value across a single generation, has been touted for its generality (see Okasha, 2006a). But this equation too can be derived from Equation (9) (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 symmetry structure of the RME to be fundamental leads to the following definition:

Definition. A population is an instance of the *Darwinian Dynamical Kind (DDK)* if the dynamics that govern its evolution respect the symmetry structure of the RME.

A process of population change is an instance of evolution by natural selection just if the population is an instance of the DDK.

4.2 Levels of selection

If undergoing evolution by natural selection means belonging to the DDK, then it is possible to provide sharp answers to questions about levels of selection. To begin with, this criterion sanctions some intuitively plausible examples. Consider the particle and group models of KGS (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 diploid genotypes. The fact that $\alpha_1 = \beta_1$ in this interpretation reflects the fact that the alleles in a genome (a group) share a

¹⁸ The haploid version of the RME is taken from (Page and Nowak, 2002, 97). The diploid version was constructed by the author using the discrete analogues of terms appearing in the continuous time RME for sexual reproduction as presented in Appendix A of the same paper.

common fate.¹⁹ With these assumptions, it is straightforward to show that Equations (1)–(2), which indicate the change through time of allele (particle) frequencies in the meta-population, are instances of the diploid RME.²⁰ It is also the case that Equations (6)–(8) are instances of the diploid RME for three distinct types.²¹ It follows trivially that Equations (1)–(2) jointly respect the symmetries of the RME, as do Equations (6)–(8). So if we take the symmetry approach seriously, both the population of particles and the population of groups in this model are in fact Darwinian evolvers. That is, both the population of alleles and the population of genotypes in a diploid population evolve by natural selection given the KGS recurrence relations and the fact that alleles in the heterozygote have equal fitnesses.

It is worth noting that the DDK 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 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 (since diploids do not reproduce with perfect fidelity). Thus, we must either treat sexually reproducing populations as evolving in somewhat less than Darwinian fashion 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 exact realizations of the fundamental symmetry structure that characterizes Darwinian evolution.

Of course, if we are to believe the DDK actually picks out all and only the instances of evolution by natural selection, it must rule out some putative cases. Here again we don’t have to look far for an example. To demonstrate that a population governed by a particular dynamics fails to satisfy the requirements for membership in the DDK, 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,²² it will suffice for our purposes to consider only the most obvious subset of the symmetries of Equation (10): fitness scaling. More explicitly, Equation (10) is invariant under transformations that change all fitnesses at a given time by a common 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.²³

Consider again the dynamics of group evolution expressed by Equations (6) – (8). Let us continue to identify the fitness of groups in terms of KGS’s π_i so that the fitnesses of the three group types are $\omega_0 = 2\beta_0$, $\omega_1 = \alpha_1 + \beta_1$, $\omega_2 = 2\alpha_2$. In general, the α ’s and β ’s may be causally independent of one another. Suppose that the ratio $\frac{\alpha_1}{\beta_1}$ is *not* constrained to be constant. Specifically, suppose the ratio varies under interventions that scale group fitnesses. 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 show that fitness scaling is no longer a symmetry of Equation (7) (see Appendix A). Put differently, unless the alleles in a group retain proportional fitnesses under intervention, the population of groups in our original example does not evolve in a process of natural selection.

¹⁹ It is sufficient for being an instance of the RME that the two allele types have identical fitnesses. It is not necessary for each allele in a group to share a common fate for this to be the case. (I thank Peter Gildenhuys for pointing this out.) In this sense, the condition of “common fate” on which Sober and Wilson (1994) have insisted is too strong.

²⁰ Specifically, they are instances of the diploid RME for two distinct types for which $Q_{000} = 1$, $Q_{010} = Q_{100} = \frac{1}{4}$, $Q_{110} = 0$, $Q_{001} = 0$, $Q_{011} = Q_{101} = \frac{1}{4}$, $Q_{111} = 1$, and $R_{00} = \alpha_2$, $R_{01} = R_{10} = \alpha_1 + \beta_1 = 2\alpha_1$, $R_{11} = \beta_0$.

²¹ Specifically, they are instances of the diploid RME for which $Q_{000} = 1$, $Q_{010} = \frac{1}{2}$, $Q_{011} = \frac{1}{2}$, $Q_{021} = 1$, $Q_{110} = \frac{1}{4}$, $Q_{111} = \frac{1}{2}$, $Q_{112} = \frac{1}{4}$, $Q_{121} = \frac{1}{2}$, $Q_{122} = \frac{1}{4}$, $Q_{222} = 1$ (with the remaining unspecified values of $Q_{ijk} = 0$), and $R_{00} = \beta_0^2$, $R_{01} = \beta_0\beta_1$, $R_{02} = \alpha_2\beta_0$, $R_{11} = \beta_1^2$, $R_{12} = \alpha_2\beta_1$, $R_{22} = \alpha_2^2$. Note that the values of the Q_{ijk} are just the probabilities one would expect for Mendelian segregation.

²² Determining the complete symmetry structure of a system of differential equations is generally quite difficult, though at least there are some standard tools (see, e.g., Oliveri, 2010). Doing so for discrete recurrence relations is even harder, and so far as I am aware, there are no systematic methods for doing so. This is, at least in a practical sense, a weakness of my dynamical kinds account.

²³ This particular symmetry is emphasized by Wagner (2010) in the context of Wright’s selection equation (a special case of the RME). In that paper, Wagner suggests a new fitness concept built from pairwise competition tests between types, a concept that assumes nothing about population growth or structure (unlike the notions of fitness as expected number of offspring or intrinsic growth rate for a type that I’ve been using). He shows that, when certain conditions are met, the resulting concept is representable on a ratio scale. Consequently, those same conditions guarantee that fitness scaling (with respect to Wright’s model and fitness as the growth rate of a type) is a dynamical symmetry. In effect, the new notion of fitness proposed by Wagner is what remains when one equates states of an evolving population that are connected by this dynamical symmetry.

It's also illuminating to ask whether a population of groups can evolve by natural selection while the population of constituent particles does not. According to the dynamical kinds account, we can answer sharply in the affirmative. As a proof by example, consider an infinite population of two kinds of groups, type 1 and type 2, comprising two kinds of particles, type A and type B. For simplicity, assume the groups are of fixed but arbitrary size and that they are homogeneous in that type 1 groups consist entirely of type A particles, and type 2 groups consist of type B particles. As a consequence the relative frequency of group types at any given time is equal to the relative frequencies of particle types. To make group fitnesses unambiguous, we'll suppose that groups directly copy themselves (by replicating their constituent particles) with an expected number of offspring groups of w_1 and w_2 for types 1 and 2, respectively. Suppose that groups evolve according to the simple haploid RME so that:

$$f_1 = w_1 f_1 / \bar{w} \quad (11)$$

$$f_2 = w_2 f_2 / \bar{w} \quad (12)$$

where $\bar{w} = w_1 f_1 + w_2 f_2$. The only unusual ingredient (which is not that unusual) we need to add to this model is a frequency-dependent particle fitness. Specifically, suppose that type A particles reproduce with an expected number of copies given by:

$$\alpha(p) = \frac{w_2(1-p)}{p(w_1 p + w_2(1-p))} \quad (13)$$

Qualitatively, each particle of type A can be expected to produce more offspring as the frequency of A's declines, and vice versa. Given this is the case, a little algebra is all it takes to demonstrate that particles evolve – in terms of particle frequencies and fitnesses only – according to the simple equation:

$$p(t+1) = 1 - \alpha p(t) \quad (14)$$

It's obvious from inspection that Equation 14 is *not* invariant under scaling of particle fitness, α . If one scales α by some positive real constant, k , at time 0 and then evolves one step in time according to Equation 14, one arrives at a state in which the frequency of A-type particles is $1 - k\alpha p(0)$ (and, of course, the particle fitness is $k\alpha$). On the other hand, if you evolve first and then scale the fitness, you end up in a state in which the A-type frequency is $1 - \alpha p(0)$. Thus, fitness-scaling is *not* a dynamical symmetry of the particle evolution. This is a simple system in which group frequencies change in a process of Darwinian evolution by natural selection, but the particles that comprise them do not.

5 Objections and extensions

I have argued that answering questions about the levels of selection means answering questions about the equivalence of dynamics—is the process of change in a population of groups of the same kind as the process in the population of individual organisms, and are either instances of the process of natural selection? I offered a general theory of dynamical kinds that carves dynamical processes into objective categories on the basis of symmetry structures, and so provides general answers to questions about dynamical equivalence. To apply this theory to questions about selection, however, I had to assume that we already know the symmetry structure characteristic of evolutionary change, and that this is captured by the symmetry structure of the RME. There are a number of obvious objections to this procedure that bear comment.

5.1 Wrong theory of evolutionary dynamics

To begin with, one might simply object that the RME is not the correct choice of theory from which to extract a fundamental symmetry structure for evolution by natural selection. Glymour (2011), for instance, has argued that the population-level approach of classical population genetics is fundamentally misguided, and that individual-based modeling offers superior predictive and explanatory power. Thus, the symmetry structure of the RME is irrelevant. I am sympathetic to this line of criticism, and I am happy to concede

that, in fact the RME is not the right theory of evolutionary dynamics to use.²⁴ But this is only an objection to the details of the analysis of the levels of selection offered above. The point is that whatever the correct theory or theories of evolutionary dynamics may be, we can apply the theory of dynamical kinds to it in order to decide at what levels of the biological hierarchy evolution by natural selection occurs. In fact, we need not tie ourselves to theoretical models at all. As I said above, we could take a data-driven approach to identifying the relevant symmetry structure by examining the symmetries of actual populations. A different choice of symmetry structure means a different definition of the DDK, but the procedure is the same. And it is the procedure for which I am arguing.

5.2 There is no single symmetry structure of evolutionary dynamics

One might instead worry that there just isn't a 'fundamental' symmetry structure that captures evolution by natural selection. Clearly there are some population-level regularities in the biological world, so the worry must be that there is such a variety of dynamics governing change through time of populations that there is no single dynamical kind that encompasses them all. Again, this may in fact be the case. However, in that event, the category 'population undergoing evolution by natural selection' is just not a natural kind. The question of levels of selection would then be ill-posed. We would instead have to ask about where in the structural hierarchy this or that more specific evolutionary process may occur. It would be a mistake to speak of levels of selection as though that were a single sort of thing. But for each sort of evolutionary process there is, the dynamical kinds approach gives clear, well-motivated answers to questions about where in the structural hierarchy such change can occur.

5.3 We don't know what the right symmetry structure of evolutionary dynamics is

Alternatively, one might think that there exists one or more kinds of evolutionary process but that we just don't know what they are. Here, the theory of dynamical kinds offers significant practical assistance. Dynamical kinds are discoverable without prior knowledge of any particular dynamics. In fact, it is possible to search for dynamical kinds without even postulating any particular dynamics. This can be done by looking directly for dynamical symmetries (Jantzen, 2014). These are often easier to discover than specific dynamical relations—one need only assess whether particular transformations commute with the time-development of the system in question (Jantzen, 2017a). Thus, not only is this worry not a problem for the dynamical kinds approach to the levels of selection, it is an invitation to test the inductive power of these putative natural kinds.

5.4 Real populations

Finally, one might object to the fact that my entire discussion has been framed in terms of infinite populations. Such populations evolve deterministically. More to the point, they exhibit exact symmetries in my sense of the term. Yet, no real populations have these features (though they approximate them to varying degrees). So no real populations meet the conditions for membership in a dynamical kind. What good is a theory of dynamical kinds that depends critically upon an idealizing assumption known to be false and unrealizable in nature? No matter how sharp the answers it gives for ideal populations, how can we take seriously a theory of kinds of process that does not seem to admit the actual, seemingly stochastic processes of real biological evolution? This worry is entirely fair given what's come before, but it's ultimately unwarranted. It is only for rhetorical expediency that I have presented the theory of dynamical kinds and the underlying notion of a dynamical symmetry in terms of exact commutation with deterministic laws of system evolution. Neither the overall process-first approach I'm advocating nor my theory of dynamical kinds in particular are limited in this way. In fact, I have already presented a generalized version of the dynamical kinds theory that

²⁴ The RME is really only a partial theory. Just like Newton's Laws of Motion are empty without a specification of one or more force laws, the RME is empty without a further specification of the ways in which fitnesses can depend on time and population states.

embraces stochastic dynamics (Jantzen, 2017a). Exploring the details would be impractical in this essay, but I encourage the concerned reader to have a look. Roughly, the idea is that we should understand dynamical symmetries not in terms of states of variables but in terms of probability distributions over those variables. At any rate, relaxing the assumption of infinite populations need not preclude the adoption of my account of dynamical kinds. And even if it did, that fact would provide no reason to reject the possibility of a theory of kinds of dynamical process.

6 Conclusion

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? What about the other way around? I suggested that evolution by natural selection is a type of dynamical process, 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 realizations, following the proposal in (Jantzen, 2014). In short, two populations whose states realize the same symmetry structure under the relevant set of transformations are said to instantiate equivalent dynamics. Viewed within this framework, to be a Darwinian evolver is to possess a certain dynamical symmetry structure. For purposes of illustration, I chose the symmetry structure of the RME as fundamental, and used this to offer a precise characterization of the Darwinian Dynamical Kind (DDK)—a category of dynamical system that may be said to undergo evolution by natural selection. I then demonstrated that it possible for populations of entities at multiple levels of the structural hierarchy to simultaneously belong (or fail to belong) to the DDK. In particular, we can clearly characterize cases in which a population of particles evolves in a process of natural selection while groups of those particles do not, as well as cases in which a population of groups belongs to the DDK but the meta-population of particles comprising the groups does not.

Whether or not the particular choice of my theory of dynamical kinds or of the symmetry structure of the RME are ultimately correct or defensible, the overall strategy is, I suggest, compelling. If evolution by natural selection is in fact a type of process, then we ought to attempt to characterize the process directly, and only secondarily worry about the properties that populations of organisms must possess to participate in such a process. In order to characterize the process in a manner that doesn't beg the question, we need a theory of kinds of process or, as I've said interchangeably, a theory of the sameness of dynamics. The analysis of Section 4.2 can be treated as a template for extracting clear and informative (and sometimes surprising) answers about the levels of selection from such a theory.

A Appendix: Violating the fitness-scaling symmetry

It is straightforward to show that fitness scaling is not a symmetry of Equations (6) – (8) unless the fitnesses α_1 and β_1 are directly proportional under all possible transformations of the system that scale group fitnesses. Each of the dynamical equations in this case is a difference equation, expressing a type frequency at time $t + 1$ in terms of type frequencies and fitnesses at time t . For ease of notation, let $\mathbf{s}(t) = \langle f_0(t), f_1(t), f_2(t), \omega_0(t), \omega_1(t), \omega_2(t) \rangle$ stand for state of the system at a time t . The variables ω_i are meant to represent the group type fitnesses. Group fitness ω_i in the discrete-time case represents the expected total number of offspring produced by a group of type i . Group fitnesses do not appear explicitly in Equations (6) – (8), but there is a simple connection between them and the particle fitnesses $\alpha_1, \alpha_2, \beta_0, \beta_1$. Since groups are assumed to be of fixed size, it must be the case that each ω_i is proportional to π_i as defined by Kerr and Godfrey-Smith (2002). Really, each group fitness is the average of the particle fitnesses weighted by the proportion of each particle type appearing in the group. In the two-particle case, this means that $\omega_0 = \beta_0$, $\omega_2 = \alpha_2$, and $\omega_1 = \frac{1}{2}(\alpha_1 + \beta_1)$. It is this last fitness relation that is problematic.

To see why, we have to consider the relevant transformations. We are interested in fitness-scaling as a set of transformations of *group-level* variables. There is one fitness scaling transformation, σ_k for every positive real-valued k . These transformations have the effect of multiplying all group fitnesses by a common factor:

$$\sigma_k(\mathbf{s}(t)) = \langle f_0(t), f_1(t), f_2(t), k\omega_0(t), k\omega_1(t), k\omega_2(t) \rangle$$

According to the definition of dynamical symmetry, σ_k is a symmetry of the group dynamics just if the system state is the same whether we apply σ_k and then use Equations (6) – (8) to evolve the system, or first evolve the system and then apply σ_k . Let Λ_i be the function mapping the state of the system at time t to the value of f_i at time $t + 1$. If σ_k is a symmetry, then it must be that for $i = 0, 1, 2$:

$$f_i(t + 1) = \Lambda_i(f_0(t), f_1(t), f_2(t), k\omega_0(t), k\omega_1(t), k\omega_2(t)) \quad (15)$$

Consider just $f_0(t+1)$. Dividing the numerator and denominator of the right-hand side of Equation (6) by $\frac{1}{2}\beta_1$ gives:

$$f_0(t+1) = \frac{(2\frac{\beta_0}{\beta_1}f_0(t) + f_1(t))^2}{\left(2\frac{\alpha_2}{\beta_1}f_2(t) + 2\frac{\beta_0}{\beta_1}f_0(t) + \frac{(\alpha_1+\beta_1)}{\beta_1}f_1(t)\right)^2} \quad (16)$$

A transformation mapping ω_0 to $k\omega_0$ is identical with one which maps β_0 to $k\beta_0$ at the particle level. The same scaling transformation must also take α_2 to $k\alpha_2$. However, the mapping from ω_1 to $k\omega_1$ does not correspond to a unique transformation of particle-fitnesses. In the most general case, we have two functions, g and h such that $g(\alpha_1) + h(\beta_1) = k(\alpha_1 + \beta_1)$. Any choice of functions satisfying this condition is equivalent to the single group fitness transformation—the transformation is infinitely degenerate from the particle perspective. However, if the transformation is to satisfy (15), then it must be the case that the coefficient of $f_1(t)$ in the denominator is constant under the scaling transformation. In other words, it must be the case that:

$$\frac{g(\alpha_1) + h(\beta_1)}{h(\beta_1)} = \frac{\alpha_1 + \beta_1}{\beta_1} \quad (17)$$

Since $g(\alpha_1) + h(\beta_1) = k(\alpha_1 + \beta_1)$, we have that $\frac{k(\alpha_1+\beta_1)}{h(\beta_1)} = \frac{\alpha_1+\beta_1}{\beta_1}$, and thus $h(\beta_1) = k\beta_1$. Likewise, since, $\frac{g(\alpha_1)+k\beta_1}{k\beta_1} = \frac{\alpha_1+\beta_1}{\beta_1}$ it must be that $g(\alpha_1) = k\alpha_1$. Thus, in order for the group fitness scaling transformation to be a symmetry, it must be the case that:

$$\frac{g(\alpha_1)}{h(\beta_1)} = \frac{\alpha_1}{\beta_1} \quad (18)$$

In plain language, the ratio of the particle fitnesses must remain fixed under all transformations of the group fitnesses. As a consequence, if the ratio of α_1 to β_1 is not fixed under all group fitness scaling transformations, then fitness scaling is not a symmetry of the group-level dynamics, and the population of groups is *not* a Darwinian evolver.

References

- Bourouina, T. and J.-P. Grandchamp (1996, December). Modeling micropumps with electrical equivalent networks. *Journal of Micromechanics and Microengineering* 6(4), 398.
- Brandon, R. (1982, January). The Levels of Selection. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 1982*, 315–323.
- Damuth, J. and I. L. Heisler (1988, October). Alternative formulations of multilevel selection. *Biology and Philosophy* 3, 407–430. 4.
- Darwin, C. (1993). *The Origin of Species* (6th ed.). New York: The Modern Library.
- Dawkins, R. (1976). *The Selfish Gene*. New York: Oxford University Press.
- Dugatkin, L. A. and H. K. Reeve (1994). Behavioural Ecology and Levels of Selection: Dissolving the Group Selection Controversy. *Advances in the Study of Behavior* 23, 101–133.
- Glymour, B. (2011). Predicting populations by modeling individuals. In J. K. Campbell, M. O'Rourke, and M. H. Slater (Eds.), *Carving nature at its joints: natural kinds in metaphysics and science*, Topics in contemporary philosophy. Cambridge, Mass: MIT Press.
- Godfrey-Smith, P. (2007). Conditions for evolution by natural selection. *The Journal of Philosophy* 104(10), 489–516.
- Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. Oxford; New York: Oxford University Press.
- Hamrin, S. F. and L. Persson (1986). Asymmetrical Competition between Age Classes as a Factor Causing Population Oscillations in an Obligate Planktivorous Fish Species. *Oikos* 47(2), 223–232.
- Hull, D. L. (1980). Individuality and Selection. *Annual Review of Ecology and Systematics* 11, 311–332. 1.
- Jantzen, B. C. (2014). Projection, symmetry, and natural kinds. *Synthese* 192(11), 3617–3646.
- Jantzen, B. C. (2017a). Dynamical Kinds and their Discovery. *Proceedings of the UAI 2016 Workshop on Causation: Foundation to Application*. arXiv: 1612.04933.
- Jantzen, B. C. (2017b). Symmetry and Causation: A General Theory of Biological Individuality. *Minnesota Studies in Philosophy of Science*. Forthcoming.
- Kerr, B. and P. Godfrey-Smith (2002). Individualist and Multi-level Perspectives on Selection in Structured Populations. *Biology and Philosophy* 17, 477–517. 4.
- Krupp, D. B. (2016, April). Causality and the Levels of Selection. *Trends in Ecology & Evolution* 31(4), 255–257.
- Lewontin, R. C. (1970). The units of selection. *Annual Review of Ecology and Systematics* 1, 1–18.
- Lloyd, E. (2012). Units and Levels of Selection. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2012 ed.).
- Lloyd, E. A., R. C. Lewontin, and M. W. Feldman (2008). The Generational Cycle of State Spaces and Adequate Genetical Representation*. *Philosophy of Science* 75(2), 140–156.
- May, R. M. and A. R. McLean (2007). *Theoretical Ecology* (3rd ed.). New York: Oxford University Press.
- Okasha, S. (2006a). *Evolution and the levels of selection*. Oxford; Oxford; New York: Clarendon Press ; Oxford University Press.
- Okasha, S. (2006b). The Levels of Selection Debate: Philosophical Issues. *Philosophy Compass* 1(1), 74–85.
- Okasha, S. (2011). Précis of Evolution and the Levels of Selection. *Philosophy and Phenomenological Research* 82(1), 212–220.
- Okasha, S. (2016, June). The Relation between Kin and Multilevel Selection: An Approach Using Causal Graphs. *The British Journal for the Philosophy of Science* 67(2), 435–470.

- Oliveri, F. (2010, April). Lie Symmetries of Differential Equations: Classical Results and Recent Contributions. *Symmetry* 2(2), 658–706.
- Page, K. M. and M. A. Nowak (2002). Unifying evolutionary dynamics. *Journal of Theoretical Biology* 219, 93–98.
- Pastor, J. (2011). *Mathematical ecology of populations and ecosystems*. John Wiley & Sons.
- Sober, E. (1993). *The nature of selection: evolutionary theory in philosophical focus*. Chicago: University of Chicago Press.
- Sober, E. and D. S. Wilson (1994). A Critical Review of Philosophical Work on the Units of Selection Problem. *Philosophy of Science* 61, 534–555. 4.
- Spirtes, P. (2009). Variable definition and causal inference. In *Proceedings of the 13th international congress of logic methodology and philosophy of science*, pp. 514–553.
- Sterelny, K. (1996). The Return of the Group. *Philosophy of Science* 63(4), 562–584.
- Vrba, E. S. and S. J. Gould (1986, April). The Hierarchical Expansion of Sorting and Selection: Sorting and Selection Cannot Be Equated. *Paleobiology* 12(2), 217–228.
- Wagner, G. P. (2010). The measurement theory of fitness. *Evolution* 64(5), 1358–1376.
- Wilson, D. S. and E. Sober (1989). Reviving the superorganism. *Journal of theoretical biology* 136, 337–356. 3.
- Wilson, R. A. (2003). Pluralism, Entwinement, and the Levels of Selection. *Philosophy of Science* 70(3), 531–552.
- Woodward, J. (2015, July). The problem of variable choice. *Synthese* 193(4), 1047–1072.

Acknowledgements I am grateful to Peter Gildenhuys and two anonymous referees for their helpful comments on earlier versions of this paper.