Symmetry and Causation:

A General Theory of Biological Individuality

Benjamin C. Jantzen

Abstract

I propose and defend a method of identifying individuals that is applicable across the biological sciences and yet sensitive to the details of particular theories. Specifically, I propose that an individual with respect to a given biological theory is an entity that instantiates the structure of a special class of transformations called the 'dynamical symmetries' of the theory. Here, a dynamical symmetry is understood roughly as a transformation of the state of a system that commutes with the increment of another system variable. This notion of individual is dependent upon the causal regularities in a particular domain of biology. However, the approach is completely general in that the same characterization of 'individual' in terms of symmetries applies across all biological domains. The metaphysical and methodological appeal of this approach to identifying individuals derives from the fact that the entities identified in this way share robust causal features and yet are causally independent of one another. To demonstrate the generality as well as the plausibility of the approach, I consider examples from evolutionary theory and ecology.

0. Overview

A motley collection of problems exercising both biologists and philosophers of biology concern the identification of biological individuals: What is an organism (Pepper & Herron, 2008)? What are the units of natural selection (Godfrey-Smith, 2009; Okasha, 2006; Sober & Wilson, 1994; Sterelny & Kitcher, 1988; Wilson & Sober, 1989)? What are species (Wheeler & Meier, 2000)? Which entities should we identify with the terms in our ecological models (Wagner & Laubichler, 2000)? Despite the apparent diversity, these questions can be seen as instances of a general, twofold problem of delineating individuals in the domain of biology. On the one hand, there is a methodological need to identify the relevant units of study when constructing generalizations or testing theories. This sort of problem is especially pressing in those disciplines like evolutionary biology and ecology that rely heavily on comparative methods (Cracraft, 1989). On the other hand, there is a question of interpretation: given a mature biological theory such as evolution by natural selection, the problem is to determine how the theory carves the world into individuals—to what particular things does the theory in fact refer? In other words, given a theory of biology that is assumed to describe the world successfully, what entities does the theory recognize as biological units? Historically, answers to each of these two types of question have been pursued separately and parochially—the practical identification of units of study has generally been ignored when considering the metaphysical question of the ontology supported by a biological theory, and questions concerning the interpretation of any one theory such as evolution by natural selection have been undertaken without regard to the interpretation of others, such as those concerned with ecological dynamics or developmental processes.

While isolated consideration of each interpretive or methodological problem of biological individuality has been productive it has not been decisive, and a generic problem cries out for a generic solution. This essay advances just such a solution. I propose an approach to identifying biological individuals that is global—the approach applies across biological theories and phenomenal domains—and that suggests an intimate connection between the individuals that populate our best interpretations of a science and those that are practical units of study. The proposed approach makes central use of dynamical symmetries, properties of the causal structures which account for the regularities amongst biological variables. Very roughly, the individuals described by a theory are those concrete biological systems that instantiate the structured collection of dynamical symmetries characteristic of that theory. The individuals so identified are independent of one another in their properties, and thus reflect genuine divisions of the world. That same independence, along with the fact that dynamical symmetries can be identified even in the absence of theory, makes these individuals salient units of empirical investigation. In other words, they are good candidates for the individuals we are seeking both for practical and interpretive purposes.

Both the way I have framed the problem of biological individuality and the approach I am offering towards its resolution are unorthodox. Even when a general problem of identifying biological individuals is acknowledged it is often assumed that we are looking for a univocal answer in terms of a collection of features a thing must have to be an individual. Hull, for instance, assumed that all individuals are "spatiotemporally localized entities that have reasonably sharp beginnings and endings in time," and are "well-integrated" (1980). He then famously argued that species names in fact refer to such individuals, rather than to intensionally defined classes. However, Hull does not tell us much about what it means to be "well-integrated" or why integration and spatiotemporal localization are the salient factors with which to divide a world of biological particulars into units. Why this principle of individuation and not some other? Even if we accept spatiotemporal localization as a necessary condition for individuality, why must the units of evolution be well-integrated in the same way as, say, the units of physiology? What reason is there to suspect this way of dividing biological particulars always reflects a genuine division in the world? In short, why think there is a unique principle of individuation in biology such that all biological domains are concerned with individuals of the same basic sort?

The same assumption of uniqueness looms large in the more recent literature on biological individuality. Ellen Clarke, for instance, remarks with reference to various candidate definitions of 'organism'—in this case treated as a synonym for 'biological individual'—that "…in order to preserve the organism definition's generality we will always need to make a choice of one candidate over another. To the extent that they are inconsistent (that they give different verdicts on the individuality of real-life cases), these candidates cannot all be right" (2010, p. 323). In motivating the need for an account of biological individuality, she cites both methodological and interpretive concerns arising from distinct theoretical problems. But there is no reason on this basis to insist that the individuals identified by evolutionary theory are coextensive with the individuals picked out by developmental science, or even that there is just one sort of evolutionary unit. Both domains might be engaged with biological individuals of very different types, and so it would be misguided to insist on a singular characterization.

The approach defended here is overtly pluralist in that it recognizes as many kinds of biological individual, carved out by as many principles of individuation, as there are stable causal generalizations in biology. I am not asserting a unique set of necessary and sufficient conditions in terms of static properties a thing must possess. Rather, I am offering a recipe for generating a principle of individuation for each dynamical generalization. Each distinct generalization, may in fact carve the world in a distinct way. There is no one set of properties or features that makes a thing an individual. Rather, there is one general way to figure out what those features are in a given domain or with respect to a given theory. In this view, it is thus an empirical fact whether the individuals described by different theories are coextensive. It is not a failure of my account of individuality if there turn out to be multiple, non-coextensive kinds of individual in biology. In fact, it is to be expected. After all, different theories may accurately describe different fault lines or joints in nature.

It is important to stress what this account of individuality is not intended to accomplish. I am not attempting a conceptual analysis with the aim of cataloguing the ways in which the term "individual" is applied in biology. Nor is it my intention to characterize what some particular set of paradigm cases of biological individuality have in common. If my account captures all or most such widely recognized cases, so much the better. But if it captures relatively few, this does not constitute failure. What matters is whether the units identified as biological individuals succeed in addressing the concerns over theory interpretation and biological methodology with which I began. If the proposal is successful in these respects, the units so identified demand our attention and recognition. If conventions of linguistic usage in biology or philosophy ultimately motivate calling these units something other than "biological individuals," so be it.

The method for individuating biological entities that I wish to defend depends upon some technical notions from a theory of natural kinds (see (Jantzen, 2014)), notions that require some effort to unpack. In order to provide this necessary technical apparatus while keeping the ultimate aim in view, the essay is structured as follows. In the next section, I make a more careful case for recognizing a general problem of individuality in biology. In Section 2, I provide a thumbnail sketch of the proposed solution. This is followed by a detailed presentation of the relevant concepts of 'dynamical symmetry' and 'symmetry structure', framed in terms appropriate for the causal regularities of biology. In Section 4, the resulting approach to identifying individuals is defended on metaphysical and methodological grounds. Finally, in Section 5 the method is applied to a pair of problems in biology. The reader who wishes to avoid technical detail while still getting a sense for what the symmetry approach offers is encouraged to read the summary of Section 2 and the examples of Section 5.

1. Troubles with biological individuals

I suggested above that biology lacks a clear, consistent method for identifying biological individuals, whether from a practical or an interpretive perspective. This deficiency is easy to miss since many biological phenomena occur at the size and time scales of everyday experience. Quite a few biological entities such as squirrels and oak trees are ordinary objects accessible to unaided human perception. It's difficult to see how there could be any serious doubt as to what the individuals are, and it's tempting to conclude that the individuals of biology are simply given to us. There a number of facts however, that belie confidence in this initial assessment. To begin with, there are many portions of the biological world whose division into individuals is controversial. Clonal stands of trees (Grant, Mitton, & Linhart, 1992), giant networks of fungal hyphae (Brasier, 1992; Smith, Bruhn, & Anderson, 1992), slime molds (Buss, 1999; Folse & Roughgarden, 2010), biofilms (Ereshefsky & Pedroso, 2013), and even ordinary flowering plants (Clarke, 2011) have all been sliced different ways by scholars looking to carve out "the" individuals. This disagreement makes even seemingly straightforward questions such as, "What's the biggest organism?" utterly intractable.

Perhaps the biggest interpretive problem in biology and philosophy concerns the individuals of evolutionary theory. A typical way to frame the problem is in terms of "which entities undergo natural selection and what it is that fits them for that role" (Lloyd, 2017, p. 19). Alternatively, as I have argued elsewhere (Jantzen, 2017b), one can productively reframe the question in terms of which biological entities instantiate a process of evolution by natural selection. Either way, the question is about how we should interpret what is taken to be a mature theory. To what things in the world does evolutionary theory refer?

However, the question of units of selection is not only a question of interpretation. Evolutionary analysis requires the identification of a population. If you don't know which entities are the individuals, you cannot determine a population, and therefore cannot identify type fitnesses. As Clarke puts it, "[i]n order to test predictions generated using modern evolutionary theory, biologists need to measure fitness. There is a lot of controversy regarding the correct understanding of the fitness concept, but no matter what interpretation of fitness you favor, measuring fitness requires being able to count individuals"(2011, p. 231). The absence of a clear notion of evolutionary individual generally stymies the application and testing of evolutionary models. These models contain terms ostensibly representing not only relative fitnesses, but such things as the frequency of variants (usually genes), migration rates, mutation rates, etc. In order to represent actual biological processes in the formalism of a model, it is necessary to map the theoretical variables to empirically accessible biological quantities. But to measure the frequency of variants within an evolving population, one has to know which things in the world are supposed to be considered members of the population. Should one count single ants or whole ant colonies? Without a clear method for individuating evolutionary units, there is no way to consistently map biological entities and properties to variables in evolutionary models, and without a coherent way to coordinate variables of models with evolutionary individuals and their properties, there is no consistent way to test or to select among them, inhibiting theoretical progress.

This difficulty of identifying evolutionary individuals and aligning them with theoretical variables is just one instance of a broader difficulty in coordinating mathematical biological models with the phenomena of biology—what Wagner and Laubichler (2000) call the "structural deficiency" of mathematical models in biology. One example they emphasize involves an experimental test of metapopulation dynamics—the theory of how populations of populations of prey interact with predators. The study in question (Murdoch, Swarbrick, Luck, Walde, & Yu, 1996) involved intervening on the distribution of scale insects (an agricultural pest) between a 'refuge' location near the center of a tree where they can remain unmolested by their principal predator, and a location amongst the branches of the tree. The experiments failed to detect the predicted refuge effect—the stabilization of the metapopulation by the presence of a refuge. Murdoch, et al. note that this result is inconclusive because it is possible that they simply misidentified the dynamically relevant populations. In other words, they may not have identified the individuals of metapopulation dynamics. This, say Wagner and Laubichler, is a common affliction of attempts to test theories of system dynamics: "We just do not have any straightforward criteria how to identify the appropriate dynamical units in a given experimental setting" (Wagner & Laubichler, 2000, p. 22). Similar considerations apply to studies of the adaptive evolution of life history which are often in conflict with experimental results. According to Wagner and Laubichler, much or most of the conflict between theory and experimental results can be attributed to a misidentification of the relevant functional units. The upshot is that a general method for identifying the relevant sort of biological individual is essential for model testing and comparison.

This methodological need goes beyond the need for clear communication or sensitivity to context. Contra Pepper and Herron (2008) and more recently Kovaka (2015), it's not enough that biologists clearly communicate their methods of individuation or adapt their questions about process to a given choice of individuals. From the perspective of theory testing, how biologists individuate must be dictated by features of the purported laws to be tested, but without begging the question. That is, it's important that we have a clear method for recognizing, e.g., the dynamically relevant populations of a metacommunity if we are to test generalizations about them and, of course, that method must not presuppose the truth of the regularities being tested. Furthermore, from the perspective of empirically driven discovery of individuals, we must have methods of individuation that carry over unequivocally from one observation or phenomenon to others. If I can't determine the sizes of my populations of aphids and ladybugs in the same way as you determine populations of hares and lynxes, we cannot pool our observations to generalize about population growth or predator-prey interactions.

As a final example, consider the problem of species. I don't mean the question of whether species are particulars or classes—I assume that in referring to a species we are referring to a certain sort of biological particular. Rather, the species problem is that of determining which collections of particular biological entities constitute the fundamental units of biological variation—the taxonomic individuals. As it happens, most biologists are content to eschew consideration of the problem and simply operate under the so-called 'Biological Species Concept.' Ubiquitous in introductory biology textbooks, this view defines 'species' as "…a sexually interbreeding or potentially interbreeding group of individuals normally separated from other species by the absence of genetic exchange" (Strickberger, 1996, p. 230). But most biologists are not attempting to catalog the diversity of life in such a way that significant patterns are drawn out, nor is every biologist in the business of explaining the historical origin and development of such patterns. Amongst the systematists—those who do make it their business to comprehend and classify the diversity of life—there is a notable lack of consensus on just what a species is. There are many competing species concepts and little agreement as to which should be favored. Aside from the Biological Species Concept, there are various versions of the 'phylogenetic species concept' (which takes species to be minimal groupings under which descent is hierarchical) (Davis, 1997), the 'monophyletic species concept' (which takes species to be maximally related groups sharing a single common ancestor) (Davis, 1997), and the ecological species concept (which views species as independently evolving lineages occupying a distinct adaptive zone) (J. S. Wilkins, 2009) to name a few exemplars. Mayden (1997) counts as many as 22 species concepts in the modern literature, and Wilkins brings the total to 26 in his more recent reviews (2006; 2009). Each of these proposals has its own advantages and disadvantages with regard to cataloging and understanding the origin of the diversity of life. For instance, there are many unsatisfactory features of the Biological Species Concept. For one, it applies only to sexually reproducing organisms. For another, it obscures variation to be found between allopatric groups. Each competing concept of species also carries with it a distinct set of theoretical convictions and empirical consequences. I won't attempt to defend a view here, but merely to emphasize the scope, longevity, and importance of the species problem. Even if species were the only sort of individual in dispute, biology as a whole could still be said to be in need of a method of identifying individuals.

From the examples considered above (and there are plenty more to be had), we can discern many instances of the two sorts of question concerning biological individuals with which I began. The first is interpretive and ontological: according to the best theories, what is the furniture of the biological world? Or perhaps more to the point, which things in the world are biological individuals? The second sort of question is eminently practical: over what empirically accessible entities should we quantify when conducting comparative studies? How do we line up the variables of our formal models with objects in the world? What should we measure or count to estimate the values of parameters in these models? It's worth stressing that these problems are not restricted to evolutionary theory. We should see the problem of biological individuation as extending across biology (and arguably beyond).[1]

---

[1] This claim is contrary to much of the current literature. (Clarke, 2013), for instance, treats the problem as specific to evolutionary biology.

2. A sketch of the solution: the causal structure of individuality

What is needed for resolving both of the problems raised above is a general method for individuating biological objects. There are some considerations that suggest the form a general solution would have to take. For one, the results of applying the method should depend sensitively on the laws or causal regularities known to govern a domain. After all, the laws are about the individuals, and so the parts of the world we isolate as individuals should reflect the content of the laws. On the other hand, if the laws are discoverable from the behavior of individuals, then we should be able to discover the individuals without directly appealing to the as-yet-unknown laws. These two considerations suggest that we ought to be looking at properties of the laws or causal regularities governing entities of a domain rather than at the causal regularities themselves. There are methodological considerations as well: in order to have any bearing on the methodological problems discussed above, the method of identifying individuals must be practical in the sense that we must have empirical access to whatever entities are selected. Finally, if the method is to give a satisfactory interpretation of biological theory, there must be compelling reason to recognize the objective existence of the individuals it identifies. Put another way, there must be a compelling reason to see the proposed approach to individuation as identifying natural divisions in the world—as 'carving nature at its joints'.

My aim in this essay is to present just such a method of identifying biological individuals. Here is the proposal in a nutshell: the causal regularities in a given biological domain (those that are captured by particular biological theories) exhibit characteristic 'symmetry structures'. Roughly, a 'dynamical symmetry' is an intervention on some of the variables of a system that commutes with the increment of another variable of the system. The composition of two dynamical symmetries is also a dynamical symmetry, and so for any set of causal regularities there is a characteristic algebra of symmetries. This algebra can be used to identify individuals: an individual of a given domain is a system which instantiates the algebra of symmetries characteristic of the causal regularities governing that domain. The idea is that biological individuals of a given type—the natural units described by this or that biological theory—are concrete particulars whose properties are governed by causal regularities that have certain properties. What individuates the biological units is not a set of static properties or a cluster of properties, but rather a feature of the way in which a collection of their concrete properties changes through time and in response to intervention. An important consequence of this approach is that all individuals of a type are guaranteed to be independent of one another in respect to the properties constitutive of the kind. It is this independence which suggests that this method of individuation really does pick out an objective division of the world. It is also of significant methodological utility—there are many reasons to select independent systems as the objects over which a theory quantifies. Of course, this is a very rough sketch. Many of the terms I've used in describing the proposal need substantial sharpening. This task is taken up in the following section.

3. The symmetry strategy

3.1 Laws, causes, and regularities in biology

By and large, biological theory little resembles the mathematical theories of physics. There may be no universal laws of biology true for all time and space, and while physical theories are typically expressed as differential or operator equations, many biological regularities cannot be represented in so precise a manner. However, both the laws of physics and the causal regularities with which biologists deal can be seen as instances of what James Woodward has called "invariant generalizations" (2003).

The class of invariant generalizations is a subclass of all generalizations which express the systematic variation of the values of one variable with the values of others. I use the term variable in the same sense as Woodward: "...variables are properties or magnitudes that, as the name implies, are capable of taking more than one value. Values (being red, having a mass of 10kg) stand to variables (color, mass) in the relationship of determinates to determinables" (Woodward, 2003, p. 39). To distinguish invariant generalizations from other generalizations, we need the concepts of intervention and invariance (sometimes called "stability"). As these are developed in detail elsewhere (Woodward, 1997; Woodward, 2000, 2003), I will only summarize the ideas here, making a few modifications helpful to our project. Briefly, "…an intervention on X (with respect to Y) is a causal process that directly changes the value of X in such a way that, if a change in the value of Y should occur, it will occur only through the change in the value of X and not in some other way" (Woodward, 2001, p. 4). To put it a little more precisely and a little more generally, suppose that we have a causal structure, G, that consists of a collection of variables and the causal relations among them. Roughly, an intervention I on a set of variables **X** in G is a manipulation of the value of the variables in **X** that satisfies the following:

M1) The change in any other variable $Y \notin \mathbf{X}$ produced by I is entailed by the causal relations of G. In other words, if the value of Y changes, then the new value y_1 of Y prescribed by G for the case $X_i = x_{i,1}$ is different from the old value y_0 prescribed by G when $X_i = x_{i,0}$.

M2) I changes another variable $Y \notin \mathbf{X}$, if at all, only through one or more variables in **X** and not directly or through some other route that does not include a variable in **X**.

M3) For any variable Y that changes as a result of I, I is not correlated with other causes of Y besides those appearing in **X** except for those falling under M1 and M2 above.

Note that the conditions as I have stated them are distinct from those appearing in, for instance, (Woodward, 2000). For one thing, I have left out one of Woodward's conditions entirely, namely that to be an intervention, the value of some variable must in fact change. Instead, I intend explicitly to include in the set of interventions on a given variable or set of variables the identity transformation—the "do nothing" intervention. Strictly speaking, this is not an intervention as Woodward defines it, but it will be helpful to include such a null operation in order to talk about the symmetries of invariant generalizations.[2]

The remaining notion we need is that of 'invariance'. "A generalization that relates changes in (or describes a correlation between) one set of variables and another is invariant if and only if it would continue to hold (or would be stable) under some intervention on variables figuring in that relationship" (Woodward, 2001, p. 5). In other words, a causal structure G is invariant if the causal relations among its variables remain intact under some range of interventions on those variables. The range of contexts in which the generalization holds—the set of interventions under which the generalization stays true—comprises the *domain* of the invariant generalization. It is robustness under intervention that distinguishes invariant generalizations from other, more ephemeral sorts. It should be clear that physical laws in this view are special cases of invariant generalizations. In particular, they are cases for which the range of interventions under which the generalization holds are unbounded. Newton's Universal Law of Gravitation for example, is supposed to hold for any

---

[2] I intend to include both 'hard' and 'soft' interventions under the general heading of 'intervention'; see, e.g., (Eberhardt & Scheines, 2007).

configuration of massive bodies. It should also be clear that, while it may be the case that biology lacks laws in some sense, it is replete with invariant generalizations ranging from the myriad causal networks of biochemistry to models of macroevolution.

3.2 Symmetries of causal structures

As I suggested above, we ought to focus not on the details of particular invariant generalizations but rather on properties of such generalizations that are themselves empirically accessible. One such property is the dynamical symmetry, which I defined previously as follows (Jantzen, 2014, p. 3633):

Definition 1 (Dynamical symmetry). Let V be a set of variables. Let s be an intervention on the variables in Int $\subset$ V. The transformation σ is a *dynamical symmetry* with respect to some index variable X $\in$ V $-$ Int if and only if s has the following property: for all x_i and x_f, the final state of the system is the same whether s is applied when X = x_i and then an intervention on X makes it such that X = x_f , or the intervention on X is applied first, changing its value from x_i to x_f , and then s is applied.

To see how this works, consider the class of 'structural casual models'. In general, we can write such a causal model with a set of structural equations of the form $x_i = f_i(\mathrm{pa}_i, \mathrm{u}_i)$, where pa_i is the set of causal parents of x_i and u_i is an error term representing noise or variables left out of the model (see, e.g. (Pearl, 2000)). For simplicity, I'll ignore the error terms, and for generality, I'll assume that f_i can be any function (often the term "structure equation model" or "SEM" is reserved for linear models). This is equivalent to assuming a complete, deterministic, and possibly nonlinear functional model. A set of structural equations is to be distinguished from a system of algebraic equations. The latter are supposed to be solved simultaneously, while structural equations are treated as independent assignments of the values on the right-hand side to the dependent variables on the left. They are independent in the sense that an intervention on x_i that sets x_i to the value x_0 has no impact on the functional form of the remaining equations (Pearl, 2000).

Given a functional causal model of this sort, we can identify a class of dynamical symmetry transformations by considering the overall state of the system. The overall state is represented simply by an ordered list of values of all variables in the model. Suppose that we have a particular structural causal model specified by a set of structural equations: $x_i = f_i(\mathrm{pa}_i)$. Interventions can be represented as a set of functions used to assign new values to some of the variables. Following Pearl (2000), I'll denote the change of the value of x_j to its new value under the transformation s by do(x_j = s(x_j)). Now consider the transformation s which results in do(x_i = s(x_i)) for all $i \in I_s$ where I_s is an index set labeling the variables of the model that are affected by the transformation. Then after s is applied to a state of the system, we have for all $j \notin I_s$, $x_j = f_j(\mathrm{s}(\mathrm{pa}_j))$. Consider a variable $x_{\mathrm{index}} \notin I_s$. Recall that a transformation of a system is a dynamical symmetry if we obtain the identical final state whether we apply s and then increment x_{index}, or first change x_{index} and then apply s.

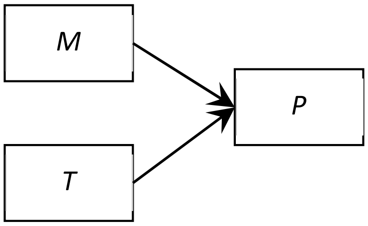

Figure 4.1

This is all very abstract; a concrete example will help to clarify. Consider the causal model indicated in Figure 4.1 and the following simple set of structural equations:

$$P = kM + h_0 \qquad (5)$$

This structural model provides a simplified picture of the relation between the internal pressure of a cell, P, and the difference between the molar concentration of solutes within and without the cell, M. The parameter h_0 indicates the pressure when $M = 0$, and k is simply a parameter capturing the effects of temperature and solute properties. Let's consider interventions on the pressure P (which can be achieved by, e.g., adjusting the hydrostatic pressure) and, since there are only two variables to choose from, we'll treat M as the index variable (it can be adjusted by, e.g., changing solute concentrations outside the cell). Now consider two sequences of interventions starting from a state $<M = 0, P = h_0>$. First, we'll adjust the solute concentration to $M = m$ which leads to the state $<m, km + h_0>$, and then increase the hydrostatic pressure by an additive constant, c. This gets us to a final state $<m, km + h_0 + c>$, which is exactly where we'd end up if we reversed the order of operations. Thus, any intervention that takes P to $P + c$ (for any real value of c) is a dynamical symmetry of the system (with respect to index variable, M). However, suppose we consider multiplicative scaling of pressure instead. If we adjust M first and then scale P by a factor of s, we'd arrive at a final state, $<m, s\ (k\ m + h_0)>$. Reversing the order of these interventions would have given us a final state of $<m, k\ m + s\ h_0>$. Multiplicative scalings of pressure are *not* dynamical symmetries of the system.

3.3 The structure of symmetry

Dynamical symmetries can be combined to produce other dynamical symmetries. If s_1 and s_2 are dynamical symmetries, then the operation of applying s_1 followed by s_2 is also a dynamical symmetry.[3] If we represent successive application of transformations as $s_1 * s_2$, then for every pair of

---

[3] Proof:
Let Λ represent the intervention that increments the index variable, V. Write $F(x)$ for the state to which intervention F takes state x. If s_1 and s_2 are dynamical symmetries with respect to V, then

(i) $\forall_x[\Lambda(s_1(x)) = s_1(\Lambda(x))]$, and dynamical symmetries we have identified a new dynamical symmetry transformation $s_3 = s_1 * s_2$.[4] This process can be repeated indefinitely—the set of dynamical symmetry transformations is closed under composition.

The behavior of a set of dynamical symmetries under composition is what I call a *symmetry structure*. To be a little more precise, a symmetry structure is a set of dynamical symmetries, S, along with a binary operator (composition) mapping elements of S×S to elements of S. As a matter of definition, every symmetry structure is presumed to contain an identity transformation e such that $e*s = s$ for any s in S. Symmetry structures are also taken to be associative in that, while the order in which transformations are applied may matter, it doesn't make a difference how they are grouped. That is, $s_1*(s_2*s_3) = (s_1*s_2)*s_3$ for all s_i in S. This notion of a symmetry structure is closely related to that of a group. A *group* in the mathematical sense is a set of operators along with a binary operation on the set that satisfies all the requirements of a symmetry structure as well as the following axiom: for every s in S, there exists an element s^{-1} such that $s*s^{-1} = s^{-1}*s = e$. Unlike groups, however, a symmetry structures may not contain an inverse for every dynamical symmetry.

An *instance* of a symmetry structure, $\Sigma = <S, *>$, is a physical system, p, that can be described via the variables on which the dynamical symmetries in S act and for which the behavior of the system under composition of the dynamical symmetries in S is that dictated by *.

(ii) $\forall_x[\Lambda(s_2(x)) = s_2(\Lambda(x))]$.

From (ii) it follows that $\forall_x[s_1(s_2(\Lambda(x))) = s_1(\Lambda(s_2(x)))]$, and from (i) we have $\forall_x[s_1(\Lambda(s_2(x))) = \Lambda(s_1(s_2(x)))]$. Thus, $\forall_x[s_1(s_2(\Lambda(x))) = \Lambda(s_1(s_2(x)))]$, or in more perspicuous notation, $\forall_x[s_1*s_2(\Lambda(x)) = \Lambda(s_1*s_2(x))]$. QED

[4] In general, dynamical symmetries do not commute. Thus s_2*s_1 may be distinct from s_1*s_2.

To make these considerations concrete, consider the structural equation model for cell pressure described above. Any system that meets the following conditions is an instance of the symmetry structure implied by that model:

(i) The system involves a pressure, *P*, and solute concentration, *M*;

(ii) any intervention s_c on *P* that changes it by an additive constant, c, is a dynamical symmetry with respect to *M*;

(iii) for any two real values c_1 and c_2, the operation of the dynamical symmetry $s_{c1 + c2}$ is equivalent to the successive application of s_{c1} and s_{c2}.

Aside from cells, one could build an artificial chamber enclosed by a semipermeable membrane that is an instance of this simple symmetry structure.

3.4 A definition of 'biological individual'

Now that we have characterized the notion of symmetry structure, we are finally in a position to provide a general definition of biological individual. Informally, the idea is to identify biological individuals with instances of the symmetry structure characteristic of a particular set of invariant generalizations. Somewhat more precisely, we can state the definition as follows:

Biological Individual (BI):

Let <S, *> be the symmetry structure of the invariant generalization governing the domain of interest. Then a set of variables V corresponds to a biological individual if and only if:

(i) The transformations of S on states of V conform to *. That is, the system of variables V is an instance of the symmetry structure <S, *>.

(ii) Removal of any element from V makes (i) false.

The second condition is only there to ensure that we do not associate irrelevant variables with an individual.

4. Justifying the strategy

4.1 Why take these 'individuals' seriously?

Before taking up some concrete examples, it is worth pausing to consider why the entities picked out by this account are plausible candidates for biological individuals. What is so appealing about a principle of individuation as abstract as the symmetry strategy? Why take such a notion seriously as a method for interpreting biological theory? The answer is that the entities marked out in this way are important from a methodological standpoint, and this methodological virtue militates in favor of taking them seriously from a metaphysical perspective.

Practically speaking, entities picked out by (BI) possess a number of virtues that make their empirical study and manipulation tractable. Chief among these is independence. The properties and evolution of each distinct instance of a symmetry structure are independent of those of all other instances of the same symmetry structure. Suppose A and B are instances of the same symmetry structure. Suppose further that they are distinct in the sense that they instantiate disjoint sets of variables (e.g. two different pressures, solute concentrations, etc.). Then the properties of A cannot depend upon the properties of B, at least with respect to those properties affected by the dynamical symmetries. If they did, then the behavior of states of A under the set of dynamical symmetries would vary depending on the state of B. Under any but the most trivial dependence on B, such variation would mean that, contrary to supposition, the dynamical symmetries on the states of A fail to constitute an instance of the relevant symmetry structure. The upshot is that the entities individuated in this way are necessarily independent of one another, and this offers the practical advantage of studying the internal dynamics of one such entity without having to account for the states of all others. If it were not possible to identify such systems that can be studied in isolation, both particle physics and biology would be intractable.

More to the point with respect to theory verification or model validation, such individuals can be identified without first knowing the details of the invariant generalization that governs their behavior. In the first place, this involves determining whether two sets of variables are in fact independent of one another, and there exists a plethora of methods for doing so. In the second place, one must determine whether two independent sets of variables belong to the same domain. That is, if they are individuals, are they individuals of the same kind? As I have shown elsewhere (Jantzen, 2017a), individuals (particular physical systems) can be sorted into the kinds that share a common symmetry structure without first learning a causal model of each individual. Consequently, this approach to individuals can remove the ambiguity biologists face, such as in the case of the metacommunity study of refuge populations described above. It is possible to verify in advance that one has identified dynamically relevant populations (the individuals that are the referents of the theory) without assuming the model expressed by the theory is correct.

It's also the case that this account satisfies most if not all of Beckett Sterner's (2015, pp. 614-615) four "epistemic tasks for individuality". The definition (BI) provides a coherent categorization of biological entities into things that are individuals (vis-à-vis a particular invariant generalization) and things that are not (although this classification does not admit of degrees). As I just argued with regard to the metacommunity example, this theory of individuation guides correct theory choice in biological models, whether evolutionary or not. Given that individuals in the sense of (BI) necessarily belong to natural kinds (see (Jantzen, 2014)), the theory tells us how and why classes of individuals of the same sort (corresponding to the same symmetry structure) are amenable to inductive generalizations. Sterner's fourth task involves explaining evolutionary transitions in individuality. It is less clear that the notion of individual captured by (BI) does this in and of itself. However, in the next section, I demonstrate its utility in clarifying units of selection, and this may offer an indirect route to satisfying Sterner's requirement. Even if it doesn't accomplish this last task, however, there is still a strong methodological case for taking (BI) seriously.

What of the ontological question? The metaphysical (and thus interpretive) virtue of the symmetry approach to individuals has also to do with independence. If in fact there are entities that persistently bear clusters of physical properties independently of one another, it would be a strange ontology that failed to recognize this distinction. That is, any account of the entities populating the world that cuts across these clusters of independent properties would be left with a curious fact to explain. On the one hand, if one were to claim that multiple clusters are instantiated in a single individual, the fact there is independence among its internal components would cry out for explanation. It would also be difficult to see why such an entity should be treated as a unity. On the other hand, if the properties belonging to each independent cluster were split across putative individuals, then the ontology would impose a strange sort of dependence amongst these individuals a priori. Since one would expect a satisfactory ontology to mirror or at least explain the empirically accessible structure of the world, neither of the alternative options is plausible. In short, if nature has joints, independence seems to be a strong indicator of where they are.

4.2 Meeting an objection

There is an obvious objection to the proposal sketched above: What not simply define 'individual' directly in terms of the instantiation of an invariant generalization? That is, why not say that individuals are particular instantiations of a set of variables governed by a given invariant generalization? It seems that all the virtues of the symmetry approach obtain for an alternate account in which a metacommunity with a refuge population is whatever obeys the equations of refuge dynamics, an individual osmotic cell is whatever obeys the SEM of Equation 5, and an evolutionary individual is whatever changes in time according to one's favorite evolutionary model.

The reason to favor symmetry is that insisting on a particular invariant generalization imposes too much. Instantiating a symmetry structure is sufficient for the sort of independence of properties described in the preceding section. Thus, insisting on a particular invariant generalization requires more than we need to in order to secure the features that make the entities identified 'individuals' in a metaphysically appealing sense. Each domain of inquiry recognizes its own set of variable types and transformations. In physics, we speak of position, momentum, etc. and transformations with respect to each of these. In evolutionary biology, we have population size, gene frequency, generation, etc. and transformations or interventions on each of these variables. A collection of entities that exhibit a particular symmetry structure—no matter what internal dynamics or causal relations they possess—are all guaranteed to be independent of one another, and thus symmetry alone is enough to carve up the world in a given domain.

Furthermore, as I suggested above, there are strong methodological reasons to focus on the symmetry rather than the full causal structure of invariant generalizations. First, symmetry is easier to discover. One can determine the symmetry structure of the states of a complex system without knowing anything about the internal causal structure. This view of individuals is therefore conducive to learning laws as well as interpreting them—one can pick out individuals on the basis of empirical symmetry characteristics without knowing the dynamics (Jantzen, 2017a). Furthermore, symmetries strongly constrain the details of an invariant generalization (Jantzen, 2014; Roy & Jantzen, in press). By identifying a collection of individuals on the basis of symmetry structure, one has already gone a long way toward determining the invariant generalization governing the behavior of that sort of individual. Thus, there are methodological as well as metaphysical reasons to favor the symmetry approach.

5. Symmetry at work: some case studies

5.1 Dynamics and the units of selection

As a first example of the kind of clarity the symmetry approach can provide in determining the referents of biological theories, consider the process of evolution. As I mentioned above, one perennial question in the philosophy of biology concerns the so-called 'units of selection'. Broadly speaking, we can put the issue this way: What biological entities evolve by natural selection? What are the units over which evolutionary theory generalizes? According to the account I've provided above, the answer to these questions depends on the symmetry structure of natural selection: if the dynamics of a changing biological entity respects the symmetry structure of natural selection, then it is what we might call an 'evolutionary individual', otherwise not. Of course, actually applying this strategy means identifying the characteristic symmetry structure of evolution by natural selection, which in turn requires us to choose which, among many candidate dynamical theories, we wish to treat as describing the dynamics of natural selection. Of course, the answers we get will depend on which theory we choose to interpret. I'll restrict attention to deterministic models of population dynamics, and set aside concerns over unrealistic assumptions such as infinite population size.

Within this class of evolutionary models, one stands out as maximally inclusive: the ‘Replicator-Mutator Equation’ or RME. This equation comes in two discrete-time forms, one for haploid (RME1) and one for diploid (RME2) organisms:[5]

$$f_i(t + 1) = (1/w)\, \Sigma_j\, q_{ji}\, w_j f_j(t) \qquad \text{(RME1)}$$

$$f_i(\mathrm{t} + 1) = (1/w)\, \Sigma_{jk}\, f_j(t) f_k(t) F_{jk} Q_{jki} \qquad \text{(RME2)}$$

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 a unit 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 (RME2), $F_{jk} = F_{kj}$ is the expected number of offspring any pair of individuals of types j and k produce together, and so $w_i = \Sigma_j f_j F_{ji}$. In both cases, $w = \Sigma_i\, w_i f_i$.

The RME is inclusive in the sense that most if not all standard deterministic models in population genetics may derived from it. In fact, the haploid version of the equation can be treated as a special case of the diploid version for which $\Sigma_k f_k\, F_{jk}\, \mathrm{Q}_{jki} = w_j\, q_{ji}$. Even the most prominent rival candidate for a universal characterization of evolutionary dynamics—the Price Equation—can be derived from (RME1). All of the models derivable from the RME inherit the same minimal symmetry structure. Thus, a symmetry possessed by the RME is possessed by all, and something that is not an individual with respect to the RME is not an individual with respect to any derived model. Thus, as in (Jantzen, 2017b), I’ll focus on the RME. While it is difficult to analytically derive the full symmetry structure of this equation, it is easy to identify a class of dynamical symmetries: the one-parameter group of fitness scaling transformations that map all w_i to $k\, w_i$, where $k > 0$ is a real-valued parameter. This

---

[5] The haploid version of the RME is taken from (Page & Nowak, 2002, p. 97) The diploid version is presented in (Jantzen, 2017b).

group is what I'll use with the symmetry approach to provide a necessary (though not sufficient) condition for evolutionary individuality.

So according to the symmetry approach, what sort of entities are candidates for evolutionary individuals and what sorts are ruled out? The first thing to note is that the units of selection must at the very least instantiate the variables occurring in (RME1) or (RME2). But, as with all models of population genetics, the only things that do so are populations. Thus, in this view, the units of selection must be populations. I want to stress this point since more often than not the units of selection are taken to be single organisms, like individual zebras or oak trees. Hull does so when he identifies 'interactors' with the objects of selection (Hull, 1980). The symmetry approach on the other hand, asserts that the units of selection—the biological individuals described by evolutionary theory—are in fact populations. Conceived in this way, the question of levels of selection is a question about which populations are evolutionary units: populations of genes? Organisms? Whole colonies? Even species?

The symmetry approach offers some concrete answers. For a detailed analysis, see (Jantzen, 2017b). Here, I'll just summarize some basic results. Consider, for instance, a population of alleles at a single locus (all the alleles in all organisms in a population of organisms). If the organisms housing the alleles are diploid, then the alleles cluster in isolated groups of two (the pairs of alleles in each cell). Each type of group (each genotype) can be treated as a distinct kind of environment. Alleles can then be assigned context-dependent fitnesses. That is, an allele of a given type will have a different fitness for each type of environment (i.e. each genotype) in which it occurs (this is the model presented in (Kerr & Godfrey-Smith, 2002)). Under these assumptions, and with the usual caveats concerning linkage equilibrium, etc., the frequency of alleles in the overall population of alleles (aggregating over all organisms) is governed by a dynamics exhibiting fitness-scaling invariance. That is, the population of alleles is (or at least may be) an evolutionary individual. If it is also the case that the alleles in each genotype have an identical fitness value, then it is also possible for the population of genotypes or the population of organisms housing them to be an evolutionary individual. However, if the fitnesses of the two types of allele in a heterozygote are not proportional to one another – that is, if the ratio of allele fitnesses is not constrained to remain fixed under scaling interventions on genotype frequencies – then the population of genotypes cannot be an evolutionary individual. Additionally, one can show that it is possible for a population of higher-level entities (e.g. organisms or colonies) to be an evolutionary individual even if the population of lower-level constituents (e.g. genes) is not, but only if the lower-level constituents combine in non-Mendelian ways.

The upshot is that the symmetry approach provides a sharp means of deciding whether or not a population is an evolutionary individual without having to assess the degree to which a population instantiates vague properties like fidelity of heredity, the dependence of fitness on intrinsic properties, and the 'smoothness' of the fitness landscape.[6] Furthermore, the symmetry approach indicates that it is possible for populations at multiple levels of the biological hierarchy to be evolutionary individuals simultaneously. Similarly, it is possible for a population of alleles to comprise an evolutionary individual in cases where groups of alleles (i.e. genomes) are not. In principle, it is even possible for the converse to obtain if inheritance is non-Mendelian. So in this case, the symmetry approach to identifying individuals brings clarity to the interpretation of a well-established biological theory.

---

[6] As is advocated, for example, in (Godfrey-Smith, 2009).

5.2 Trophic species

As a second example, I'd like to consider a case in which neither the invariant generalization pertaining to the domain of interest, nor the relevant symmetry structure is known. This case concerns the trophic structure of ecological communities. In part, the flow of energy and resources in a community may be characterized by a 'food web', a network indicating which types of organism eat which others. A common practice in the study of food webs is to aggregate taxonomic species into single units called 'trophic species'. A trophic species "…is a collection of organisms that have the same diets and the same predators" (Cohen, Briand, & Newman, 1990). The use of trophic species amounts to positing the existence of a particular symmetry structure.

While it would take us too far astray to canvas the debates over the use of trophic species in the ecological literature, the symmetry approach to identifying individuals allows us to draw some quick conclusions concerning the plausibility of this construct. For the moment, we'll assume that the only variables associated with the dynamics of a food web are population sizes. More precisely, to each food web there corresponds a structural causal model in which the only variables are population sizes for each type of organism related in the web. Of course, one might be interested in studying more than population dynamics, and even in population dynamics one might be interested in more than simply population sizes. Also, the traditional theoretical interest in food webs concerns regularities in their structural features across ecosystems: regularities in the links per node in the web, regularities in the number of top predators, etc. However, the links in the web have to do with predator-prey interactions, and these essentially concern the dynamics of interacting populations. We thus do little violence to the discipline by focusing on population dynamics.

According to the account of individuality I've defended, the individuals of community population dynamics – call them community individuals – are systems that instantiate a collection of variables referring to population sizes and that constitute a representation of some as yet unknown symmetry structure. From this perspective, the supposition that aggregates of taxonomic species can be lumped together into a single trophic species amounts to the claim that the dynamics governing community individuals are insensitive to the proportion of various taxonomic species within a given trophic species. It follows that those dynamical symmetries of community individuals that change the size of populations of a trophic species – whatever those dynamical symmetries may be – must not depend upon how the change in size is accomplished. In other words, if $s_1, s_2, \ldots, s_n$ represent the population sizes of n taxonomic species aggregated into a single trophic species, then each transformation from $s_1, s_2, \ldots, s_n$ to $s_1', s_2', \ldots, s_n'$ such that $\Sigma_i s_i = \Sigma_i s_i'$ must also be a dynamical symmetry for the community individuals of ordinary taxonomic species. But for any empirically plausible theory of population dynamics, this assumption in turn requires that the intrinsic growth rate of one population s_i is equal to the intrinsic growth rate for all other s_j. This is a rather dubious hypothesis. To sum up, supposing meta-populations of trophic species to be individuals of community dynamics in the same sense as meta-populations of ordinary taxonomic species amounts to positing a particular dynamical symmetry of the underlying invariant generalizations. This symmetry in turn imposes implausible constraints on the invariant generalization, and so we are led to reject trophic species as viable components of community individuals. This conclusion was reached without knowing either the invariant generalization governing trophic interactions or which collections of populations constitute trophic individuals. The symmetry characterization of individuals alone is sufficient to detect the problem.

6. Conclusions and caveats

According to the view laid out above, biological individuals are entities bearing a domain-specific set of biological properties and whose states evolve or are otherwise determined independent of one another. The individuals in this sense are delineated by appeal to the symmetry structure of the invariant generalization governing the relevant domain of biology. In particular, an individual is an instance of the symmetry structure of an invariant generalization. The independence essential to these individuals makes the account metaphysically appealing as it suggests that we have identified genuine divisions in the world. Methodologically, it selects precisely those units on which it is possible to experiment without worry over confounding factors. Perhaps more importantly, symmetry structures can be discovered independent of dynamics, and so it is possible to learn what the individuals in a given domain are even without knowing the relevant invariant generalizations. Symmetry can even be used to constrain the content of the invariant generalizations in a domain.

As I warned at the outset, this account of individuals is pluralistic. There is not one single kind of biological individual, but rather there could be as many types of biological individual as there are invariant generalizations in biology, though it is likely that at least some types are coextensive. One might interpret this situation in terms of 'degrees of individuality'. The idea would be that the more classes of individual a given entity falls into, the greater the degree of its individuality. Robert Nozick defends a view like this concerning 'objectivity', a view that also appeals to the symmetry properties of laws (Nozick, 2001). However, I see no reason to treat individuality as a variable property in this way. While it is of theoretical interest that categories of individual overlap, it is unlikely that there exist any entities that are individuals in all biological senses. Thus, one end of the supposed spectrum of individuality is empty. Furthermore, the more invariant generalizations we consider the more the characteristics of individuality conflict. It seems that a sort of pluralism is the least problematic resolution—it's best to embrace a diversity of biological individuals belonging to partially overlapping kinds.[7]

At the same time, I wish to stress that the account is not constructivist—our theories do not create or construct the individuals, but rather refer to them. As I argued in Section 4, these individuals constitute natural divisions in the world; the independence of properties is an objective feature of the world and so too are the units carved out by symmetry structures. Theories are descriptions. The symmetry approach to individuals merely identifies the concrete objects, if any, to which the theory may be said to refer. Where we exercise choice is in the selection of which properties to consider. But once such a choice is made, both symmetry structures and the individuals of my account are objective features of the world, knowable even in the absence of a relevant theory.

Furthermore, the symmetry approach is completely general. I cast (BI) in terms of the invariant generalizations that apply in biology because it was in biology that we are met with a host of puzzles concerning individuals. However, there is no reason that the same approach could not be taken in chemistry, geology, meteorology, psychology, etc. Every science has its invariant generalizations, and so the symmetry approach to individuation is applicable to every science.

It is also worth mentioning a complication I have entirely ignored: the invariant generalizations of biology (and other sciences) tend to be approximate or probabilistic. I have been assuming that they can be characterized in terms of exact symmetries. These facts are in tension with one another, but ought not to be of concern. Everything I have said about biological individuals can and should be

---

[7] This is similar to the metaphysical position advocated by C. Kenneth Waters (1998), though his "theoretical kinds" are somewhat different from the kinds of individual delineated by the symmetry approach.

developed in terms of dynamical symmetries acting on *distributions* of variable values rather than on variables themselves. And as it happens, everything I said above can be extended to the probabilistic case via the following modified definition:

Definition 2 (Dynamical symmetry). Let V be a set of random variables. Let s be an intervention on the variables in $\text{Int} \subset V$. The transformation s is a dynamical symmetry with respect to some index variable $X \in V - \text{Int}$ if and only if s has the following property: for all probability distributions f and g, the final probability distribution over V is the same whether s is applied when the distribution of X is given by $p_x(x) = f(x)$ and then an intervention on X makes it such that $p_x(x) = g(x)$, or the intervention on X is applied first, changing its distribution from f(x) to g(x), and then s is applied.

Unsurprisingly, dynamical symmetries defined in this way are a lot more complex to work with, and have been ignored purely out of pragmatic presentational concerns.

Finally, it should be emphasized that the symmetry approach to individuals in biology is not merely a tool for post hoc interpretation. Given that symmetry structures can be discovered in advance of dynamics, there are potentially useful roles the symmetry approach can play in developing biological disciplines such as metagenomics. It also stands to make headway with some old concerns like the problem of species.[8] These applications, however, will have to be left for future work.

---

[8] Interestingly, Ghiselin's "Radical Solution to the Species Problem" (1974) employs something close to a characterization of species in terms of symmetry structure, at least as Reed (1979) interprets the proposal.

Works Cited

Brasier, C. (1992). A champion thallus. *Nature (London), 356*(6368), 382-383. doi:10.1038/356382a0

Buss, L. W. (1999). Slime molds, ascidians, and the utility of evolutionary theory. *Proceedings of the National Academy of Sciences, 96*(16), 8801-8803. doi:10.1073/pnas.96.16.8801

Clarke, E. (2010). The Problem of Biological Individuality. *Biological theory, 5*(4), 312-325. doi:10.1162/BIOT_a_00068

Clarke, E. (2011). Plant individuality and multilevel selection theory. In B. Calcott, K. Sterelny, & E. Szathmáry (Eds.), *The major transitions in evolution revisited* (pp. 238-250). Cambridge, Mass.: MIT Press.

Clarke, E. (2013). The multiple realizability of biological individuals. *The Journal of Philosophy, 110*(8).

Cohen, J. E., Briand, F., & Newman, C. M. (1990). *Community food webs : data and theory*. Berlin; New York: Springer-Verlag.

Cracraft, J. (1989). Speciation and its ontology: the empirical consequences of alternative species concepts for understanding patterns and processes of differentiation. In D. Otte & J. A. Endler (Eds.), *Speciation and its consequences* (pp. 28-59). Sunderland, England: Sinauer Associates.

Davis, J. I. (1997). Evolution, Evidence, and the Role of Species Concepts in Phylogenetics. *Systematic Botany, 22*(2), 373-403.

Eberhardt, F., & Scheines, R. (2007). Interventions and Causal Inference. *Philosophy of Science, 74*(5), 981-995.

Ereshefsky, M., & Pedroso, M. (2013). Biological individuality: the case of biofilms. *Biology & Philosophy, 28*(2), 331-349. doi:10.1007/s10539-012-9340-4

Folse, Henri J., & Roughgarden, J. (2010). What is an Individual Organism? A Multilevel Selection Perspective. *The Quarterly Review of Biology, 85*(4), 447-472. doi:10.1086/656905

Ghiselin, M. T. (1974). A Radical Solution to the Species Problem. *Systematic Zoology, 23*(4), 536-544.

Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. Oxford; New York: Oxford University Press.

Grant, M. C., Mitton, J. B., & Linhart, Y. B. (1992). Even larger organisms. *Nature (London), 360*(6401), 216-216.

Hull, D. L. (1980). Individuality and Selection. *Annual Review of Ecology and Systematics, 11*(1), 311-332. doi:doi:10.1146/annurev.es.11.110180.001523

Jantzen, B. C. (2014). Projection, symmetry, and natural kinds. *Synthese, 192*(11), 3617-3646. doi:10.1007/s11229-014-0637-5

Jantzen, B. C. (2017a). Dynamical Kinds and their Discovery. *Proceedings of the UAI 2016 Workshop on Causation: Foundation to Application*.

Jantzen, B. C. (2017b). Kinds of process and the levels of selection. *Synthese*, 1-27. doi:10.1007/s11229-017-1546-1

Kerr, B., & Godfrey-Smith, P. (2002). Individualist and Multi-level Perspectives on Selection in Structured Populations. *Biology and Philosophy, 17*(4), 477-517. doi:10.1023/a:1020504900646

Kovaka, K. (2015). Biological Individuality and Scientific Practice. *Philosophy of Science, 82*(5), 1092-1103. doi:10.1086/683443

Lloyd, E. A. (2017). Multilevel Selection and Units of Selection Up and Down the Biological Hierarchy. In R. Joyce (Ed.), *The Routledge Handbook of Evolution and Philosophy* (pp. 19-34). New York: Routledge.

Mayden, R. L. (1997). A hierarchy of species concepts: the denoument in the saga of the species problem. In M. A. Claridge, H. A. Dawah, & M. R. Wilson (Eds.), *Species: the units of diversity* (pp. 381-424). London: Chapman & Hall.

Murdoch, W. W., Swarbrick, S. L., Luck, R. F., Walde, S., & Yu, D. S. (1996). Refuge Dynamics and Metapopulation Dynamics: An Experimental Test. *The American Naturalist, 147*(3), 424-444.

Nozick, R. (2001). *Invariances : the structure of the objective world.* Cambridge, Mass.: Belknap Press of Harvard University Press.

Okasha, S. (2006). *Evolution and the levels of selection.* Oxford; Oxford; New York: Clarendon Press ; Oxford University Press.

Page, K. M., & Nowak, M. A. (2002). Unifying evolutionary dynamics. *Journal of theoretical biology, 219*, 93-98.

Pearl, J. (2000). *Causality : models, reasoning, and inference.* Cambridge, U.K.; New York: Cambridge University Press.

Pepper, J. W., & Herron, M. D. (2008). Does biology need an organism concept? *Biological reviews of the Cambridge Philosophical Society, 83*(4), 621.

Reed, E. S. (1979). The Role of Symmetry in Ghiselin's "Radical Solution to the Species Problem". *Systematic Zoology, 28*(1), 71-78.

Roy, S., & Jantzen, B. C. (in press). Detecting causality using symmetry transformations. *Chaos.*

Smith, M. L., Bruhn, J. N., & Anderson, J. B. (1992). The fungus Armillaria bulbosa is among the largest and oldest living organisms. *Nature (London), 356*(6368), 428-431. doi:10.1038/356428a0

Sober, E., & Wilson, D. S. (1994). A Critical Review of Philosophical Work on the Units of Selection Problem. *Philosophy of Science, 61*(4), 534-555.

Sterelny, K., & Kitcher, P. (1988). The Return of the Gene. *The Journal of Philosophy, 85*(7), 339-361.

Sterner, B. (2015). Pathways to pluralism about biological individuality. *Biology & Philosophy, 30*(5), 609-628. doi:10.1007/s10539-015-9494-y

Strickberger, M. W. (1996). *Evolution* (2nd ed.). Boston: Jones and Bartlett Publishers.

Wagner, G., & Laubichler, M. (2000). Character identification in evolutionary biology: The role of the organism. *Theory in Biosciences, 119*(1), 20-40. doi:10.1007/s12064-000-0003-7

Waters, C. K. (1998). Causal Regularities in the Biological World of Contingent Distributions. *Biology and Philosophy, 13*(1), 5-36. doi:10.1023/a:1006572017907

Wheeler, Q., & Meier, R. (2000). *Species concepts and phylogenetic theory : a debate.* New York: Columbia University Press.

Wilkins, J. (2006). Species, kinds, and evolution. *Reports of the National Center for Science Education, 26*(4), 36-45.

Wilkins, J. S. (2009). *Species: a history of the idea.* Berkeley: University of California Press.

Wilson, D. S., & Sober, E. (1989). Reviving the superorganism. *Journal of theoretical biology, 136*(3), 337-356. doi:10.1016/s0022-5193(89)80169-9

Woodward, J. (1997). Explanation, Invariance, and Intervention. *Philosophy of Science, 64*, S26-S41.

Woodward, J. (2000). Explanation and Invariance in the Special Sciences. *The British journal for the philosophy of science, 51*(2), 197-254.

Woodward, J. (2001). Law and Explanation in Biology: Invariance Is the Kind of Stability That Matters. *Philosophy of Science, 68*(1), 1-20.

Woodward, J. (2003). *Making things happen : a theory of causal explanation.* New York: Oxford University Press.