

Defining Paradigm Darwinian Populations

John Matthewson*†

This article presents an account of the biological populations that can undergo paradigmatic natural selection. I argue for, and develop Peter Godfrey-Smith's claim that reproductive competition is a core attribute of such populations. However, as Godfrey-Smith notes, it is not the only important attribute. I suggest what the missing element is, co-opting elements of Alan Templeton's notion of exchangeability. The final framework is then compared to two recent discussions regarding biological populations proposed by Roberta Millstein and Jacob Stegenga.

1. Introduction. Until relatively recently, the concept of a biological population had not received a great deal of philosophical attention. This situation has changed in the last few years, and the literature now contains a number of different approaches to the question of how to define these populations. For example, some authors have attempted to develop a quite general characterization of populations in biology (e.g., Millstein 2006, 2010), while others have argued that there is no such singular characterization to be given (e.g., Stegenga 2014). This article takes a slightly different approach to both of these positions. Following work by Peter Godfrey-Smith, I develop a set of criteria that can be used to delineate the populations that are able to undergo *paradigmatic natural selection*. It might well be that there is no single account of biological populations to be had in general, but it is still plausible that we can provide an analysis of this specific type of population. I begin by

*To contact the author, please write to: Department of Philosophy, Massey University, Auckland, New Zealand; e-mail: j.matthewson@massey.ac.nz.

†Thanks to audiences at PBDB4, ISHPSSB 2011, AAP 2011, and Victoria University of Wellington for their input. I also acknowledge the help and support of Pierrick Bourrat, Patrick Forber, David Gilbert, Alan Hájek, Claire Matthewson, Chris Pincock, Jacob Stegenga, Kim Sterelny, and John Wilkins. Three anonymous reviewers for this journal also contributed extremely useful and generous comments. Finally, a great many thanks to Peter Godfrey-Smith and Roberta Millstein for their extensive help on earlier drafts. None of the above people necessarily agree with the ideas expressed here, and any errors are mine.

Philosophy of Science, 82 (April 2015) pp. 178–197. 0031-8248/2015/8202-0002\$10.00
Copyright 2015 by the Philosophy of Science Association. All rights reserved.

discussing populations that can undergo natural selection in a more minimal sense, and build from there.

Standard discussions of Darwinian evolution state that if a population exhibits phenotypic variation, fitness differences, and heredity, then natural selection will occur within that population. However, it is quickly apparent that more needs to be said regarding what the term “population” means in this context, because it is not the case that any group of organisms that meets these three requirements will undergo natural selection. For example, the “population” of two walruses, a dandelion, and a single bacterium exhibits phenotypic variation, fitness differences, and heredity, but this is not a collection of individuals that can undergo natural selection, because they simply do not form the right kind of group.

Traditionally, accounts of populations that may undergo natural selection have specified that the members of such populations must be close together in space and time and/or that they must all be of the same species.¹ On the surface this looks entirely reasonable. For a start, it identifies a seemingly obvious flaw with the gerrymandered group above: walruses, dandelions, and bacteria belong to different species. However, invoking species membership and proximity here is misguided. Even if it turned out that the groups that undergo natural selection are always spatiotemporally clustered or composed of only one species, it would be a mistake to define them as such. Most crucially, such an analysis would not tell us what it is about having members that are all the same species or close to one another that enables a population to undergo natural selection. It presumably is not “conspecificity” or “proximity” per se; rather, it will be some intragroup relations that either entail or are enabled by these attributes.

Additionally, as Peter Godfrey-Smith (2009) points out, our analysis must be able to extend beyond that of strict or clear species membership, since natural selection occurs in cases where applying species concepts is problematic. For example, delineating species boundaries is a difficult issue when we consider microbial organisms and many plants (see, e.g., Templeton 1989; O’Malley and Dupre 2007). For these and other reasons, many commentators have turned from conspecificity and proximity to the idea that what really matters are the *causal connections* between a population’s members. So the important project is therefore to discover which causal connections really matter.

2. Darwinian Populations and Reproductive Competition. In his 2009 book, Godfrey-Smith suggests two kinds of causal linkage that might delineate the populations that can undergo natural selection (what he calls the “Darwinian populations”): interbreeding and reproductive competition. The

1. For a selection of such accounts, see Millstein (2010, 61).

first of these is essentially the standard notion underlying the so-called biological species concept, while the second is less familiar. Godfrey-Smith gives an intuitive, first-pass picture of reproductive competition in a few ways. For example, this type of competition occurs in a population when “there is a causal dependence between how many offspring each individual has” and when “a slot I fill in the next generation is a slot you do not fill” (2009, 51). He symbolizes reproductive competition with a variable: the Greek letter α , where an α of 1 represents the maximum level of such competition.

Godfrey-Smith is clear that these two criteria are probably not sufficient for a group to form a potential Darwinian population, and interbreeding is obviously not a necessary condition, as some Darwinian populations contain organisms that reproduce asexually. What about reproductive competition? The assertion that this property is a requirement for natural selection to occur in general would be controversial, but here Godfrey-Smith makes a slightly different claim, “suggest[ing] that paradigm cases of evolution by natural selection occur in populations where α is in the vicinity of one” (2009, 52). So according to Godfrey-Smith, high α is important, but not sufficient, for a group to be a “paradigm” Darwinian population.

This marks a distinction between Darwinian populations in general and the paradigm Darwinian populations in particular. Godfrey-Smith’s approach here is to view instances of natural selection as forming a spectrum with quite “minimal” cases of natural selection at one end and much richer cases at the other. To illustrate this idea, Godfrey-Smith presents an example from Richard Lewontin (1970), where two strains of bacteria are growing in an excess of nutrient broth with no constraints on population growth. One strain reproduces more rapidly than the other, which means that over time the faster-reproducing type becomes more numerous than the slower type. In this case, it appears as though natural selection of some kind can occur even when the number of offspring of one phenotype is not at all linked to the number of offspring of the other phenotype.

However, it seems clear that Lewontin’s example represents a very minimal case of natural selection; a case importantly different to the sort of circumstances that lead to the development of complex morphological and behavioral adaptations. As I see it, there are (at least) two ways in which the Lewontin thought experiment characterizes a rather “anaemic” case of natural selection. First, both strains are reproducing as fast as they are able; it just so happens that one has a faster maximum reproductive rate than the other. As pointed out by Lennox and Wilson (1994), this stretches the notion of the bacteria undergoing selection rather far and essentially takes environmental influences out of the picture.² Perhaps we might sometimes be interested in cases like this, but a situation in which the environment plays no

2. Thanks to Roberta Millstein for bringing this argument to my attention.

role in a population's evolutionary dynamics does not seem a particularly exemplary instance of "natural selection."

Note, though, that there is a second way in which this example falls short of Godfrey-Smith's criteria. The environment can exert a selective influence on a group of individuals, while that group still lacks any internal causal connections with respect to reproductive outcomes. For example, the introduction of an antibiotic to Lewontin's broth might affect one bacterial strain more than the other. Now selective pressure from the environment may alter bacteria numbers over time, but as long as there remains plenty of food and space for all concerned, we will still not see the kind of interconnection that Godfrey-Smith describes. Strain A might deal with the antibiotic better than strain B, but if strain B's fate is completely unconnected to that of strain A, this is still not a case of reproductive competition.

Why might we be particularly interested in groups that exhibit causal interactions with respect to reproduction, and why must they specifically be competitive interactions? One reason is that these interactions influence the role that natural selection can play when explaining the presence of certain traits within a population. In paradigm cases of natural selection, the process of selection itself is an important part of such explanations, and for this to be the case, we require more than just a shared selective environment. This is because at its core, natural selection is a process of culling unfavorable types, which in turn makes it difficult to see how it can explain the presence of favorable types. Rather, the reason a favorable type is present in a current population can be explained by a mutation in some ancestor, followed by inheritance of that mutation when selection has not occurred.³

One can attempt to avoid this conclusion by invoking the presence of strong reproductive competition. The issues here are complex, but the basic idea is that in order for natural selection to explain the presence of traits, there must be cumulative selection against a background of competition over reproductive opportunities. For example, Godfrey-Smith argues that high α is important for "origin explanations": cases in which natural selection is responsible for the appearance of a novel phenotype in a population, rather than just the removal of certain existing phenotypes. I will not go into his discussion here for reasons of space, but I will mention a different argument for a similar conclusion from Bence Nanay (2005, 2010).

Nanay claims the culling of unfavorable types can explain why a particular trait is found in a population, but only under certain circumstances. Roughly put, if only a certain number of individuals in each generation can survive to reproductive age, then the deaths of some individuals in a population will play an explanatory role regarding the survival of others in that population. In the presence of reproductive competition, selection against the traits of one

3. For examples of this controversy, see Neander (1995) and Sober (1995).

of these individuals will be an important part of the explanation for the presence of traits of the other individuals in this and the next generation. Put another way, if there can only be a certain number of offspring in the next generation, and something bad happens to you because you have trait t instead of trait T like I do, this leaves more space for my T -bearing offspring. In that case, selection against you helps explain my reproductive success and thereby the presence of T in the population.

In truly paradigmatic cases of natural selection, then—cases where natural selection can feature centrally in our explanations of traits within the population—we require not just a shared selective environment, but also the presence of strong reproductive competition. As we will see, the project of defining paradigm Darwinian populations requires more than just providing an account of this reproductive competition, but I now turn to that initial task.

3. α as a Zero-Sum Game. The core notion invoked in Nanay's argument and expressed in the quotations from Godfrey-Smith above appears to be that α is high when reproduction within the population is a zero-sum game. That is, if someone in the population does better, someone must equally do worse.⁴ Intuitively this feels as though we are capturing the basic idea, but the zero-sum interpretation of α turns out to be too weak. This can be illustrated with a simple example.

Start with two paradigm Darwinian populations that have no direct connections between them with respect to reproductive competition. For example, they may be located near one another but have no ecological interactions, or they may be spatiotemporally separated. Now consider these two populations as subpopulations of a single larger population. This further population, made up of two unrelated subpopulations, is of course not very well interconnected with respect to reproductive competition, and so any reasonable account of α should register that fact. However, as the zero-sum interpretation of α stands, this gerrymandered population has a maximal level of α .

As shown in figure 1, we will call the subpopulations A and B and the population under consideration (i.e., the union of A and B) C. Choose any member of C, increase its number of offspring by one, and we are guar-

4. Another way to characterize this approach is to treat reproductive competition as a function of the relationship between two quantities commonly employed in population biology: *population size* N and the *carrying capacity* K . Since K specifies the maximum number of organisms that can be supported by the available resources, reproduction will approach a zero-sum game as N approaches K . Unfortunately, this cannot be the (whole) story, because without a prior understanding of which groups can be the bearers of N and K , this interpretation will also fall afoul of the gerrymandering problem discussed in this section.




Figure 1. Counterexample to the idea that a high α (only) means that reproduction is a zero-sum game. A and B are both paradigm Darwinian populations that together make up a further population C. However, these two subpopulations do not interact with respect to reproductive competition.

anted that some member of C will lose one of its offspring. This is because populations A and B are paradigm Darwinian populations, so any member in A who gains an offspring will cause a member in A to lose an offspring, and similarly for members of B. Since any member of C must be a member of either A or B, this means if any member of C gains an offspring, some member of C must lose an offspring. Therefore C has an α of 1 according to the current framework, even though C clearly does not form the right kind of group to count as a paradigm Darwinian population.

This shows that the zero-sum definition of α simply fails to capture what it is intended to capture: the extent to which the population is interconnected with respect to reproductive competition. We need a different account of α that actually picks out the property of interest.

4. α as a Measure of Connectedness. Luckily, the above example makes what is required very clear. It is not enough simply that some member of the population is adversely affected by an increase in the number of offspring. Rather, we are interested in the extent to which the population as a whole is linked, such that the reproductive prospects of all its members are affected by an increase in offspring somewhere within that population.

We can get some grasp of this idea if we think about the population as represented by a graph. The members of the population can be depicted as a

set of nodes, and an edge between two nodes will indicate a causal interaction with respect to reproductive output between the population members represented by those nodes.⁵

The intention here is to use the graph-theoretic framework as a conceptual tool, rather than a rigorous method of quantification. However, with this in mind, we can illustrate the general idea with a simple approach. Given Godfrey-Smith's stipulation that the maximal level of α is 1, the extent to which each member is connected to the rest of the population can now be simply quantified by considering the number of edges connected to each of the nodes. First we calculate the ratio of the number of edges connected to each node compared to the number of edges it would have if the graph were complete (i.e., where every member is connected to every other member). This can then be used to assess the average connectedness of the population's members, giving us an overall picture of the causal interconnections within the population. This means a population will have an α of 1 when all members are directly causally linked to one another with respect to reproductive competition.

Now consider population C discussed in section 3 and depicted once again in figure 2, made of two paradigm Darwinian populations that are not connected with respect to reproductive competition. We can see that this putative population fares relatively poorly, as any individual will only be connected to half of the population. The ratio of this to the optimum is therefore $0.5N / N = 0.5$.⁶ Since the chosen individual was arbitrary, the average value will be 0.5, so the α of this relatively disconnected population is now 0.5. Using this simple method, the value of α will increase in cases where one well-connected subpopulation comprises more of the whole and as connections between the groups are introduced.

4.1. Complications. This new interpretation of α appears to capture much of what we want from our account. However, it is not yet general enough. If an edge between nodes in the graph only indicates that there is some kind of causal connection between group members with respect to reproductive success, the current account will incorrectly classify certain groups as potential paradigm Darwinian populations.

5. Here I acknowledge extensive help from David Gilbert when attempting to develop a framework depicting these populations as graphs. Although much of the material does not appear in the current article, it was very important in establishing the feasibility of the account.

6. I have included the reflexive edge because each further offspring will place available resources under greater strain, even for the parent of that offspring. I am not certain that we should include this particular reflexive relation, but in any case the decision whether to include the reflexive edge will only have a small effect on our results in the majority of examples (i.e., when N is substantially greater than 1).




Figure 2. The gerrymandered population C as a graph. Nodes are depicted as circles, edges as lines. Here, we see just the edges associated with one individual. Although this individual is connected to half of the population (including itself), it is not connected to the other half. Missing edges indicate the population is less interconnected as a whole.

For example, consider a group of predators and prey. Predators can have a causal influence on the reproductive success of other predators (when they compete for the same resources) and also on their prey (since the presence of predators will make it more difficult for prey to produce surviving offspring), while prey have an effect on other prey and also on the predators (since more prey means that there are more resources available for those predators). This means every member of the group may have some causal influence on the reproductive success of every member of that group. Or consider a population made up of two subgroups with only partial and mild overlap of the resources they require, as opposed to a population composed of conspecifics under intense resource pressure. If the edges in our graphs only register qualitative effects, a graph that represents mild competition over a few resources (as in the former case) will be just the same as a graph where all members fully compete (as in the latter case).

The issues here are pretty obvious. First, not all of the interactions of interest will be symmetrical. Second, some interactions can have a positive effect on reproductive output, rather than a negative one. Finally, the kinds of interactions we are concerned with are not binary but come in degrees. Luckily, again, each of the refinements required to address these issues are similarly apparent. First, the graph must be a *directed* graph. This means that an edge from one individual to another is not the same as an edge in the opposite direction. Second, the edges need to differentiate between positive

and negative influences on reproduction. We can do this by introducing polarity to the connections. Since α is intended to measure the level of strong competition, any causal interaction that improves the chances of reproductive success for members of the population needs to be recognized as something else again. This is not to say that positive interactions are unimportant or may not be a relevant feature when determining biological populations of some kind or other. However, they are not the relations we are concerned with for the purposes of identifying the paradigm Darwinian populations.⁷ Third, the edges will need to be weighted according to the strength of the interactions they represent.

Unfortunately all of the above substantially complicates how we can assess the interconnectedness of the population. At this point the details of such considerations would outstrip the level of precision we are dealing with (not to mention my technical resources) and take us well beyond the goals of the article.⁸ In any case, even assuming our graph-based account of α has the expressive resources to differentiate most cases appropriately, important limitations remain.

4.2. Limitations of the Causal Connection Account. Consider once again a population made up of two subpopulations that are themselves paradigm Darwinian populations. This time, it is also made explicit that the subpopulations are different species from one other. Furthermore, these two groups are in direct competition over an essential resource, and the population as a whole is operating at the limit of that resource. This means any further offspring produced by a member of one of the subpopulations will have a negative effect on every member in that subpopulation and also on every member of the other subpopulation to (almost) the same extent. In such a case, our graph will be (almost) perfectly connected, and the α of the population will therefore be (almost) maximal. So even under our most sophisticated account, a population of two different species with high resource competition can exhibit near perfect levels of α .

This is a revealing case, I think, for two reasons. The first is that if we look at the internal interactions of this population in a particularly abstract way, it starts to become less clear why we should not consider this a group

7. It was pointed out to me by Cailin O'Connor that negative and positive interactions could both be considered separately. This does seem a promising extension of the framework. I say a bit more on this later in sec. 6, but how these positive and negative relations interact with one another will depend on details specific to particular cases. For now, we will just note that the negative interactions are the ones most relevant for determining α . Many thanks also to an anonymous referee for pushing me further on this point.

8. Here I acknowledge the generous help of Aleksandar Ignjatovic for discussing these issues and the tractability of this type of analysis with me.

that could undergo natural selection. Here is a population made up of two types of individuals whose reproductive output is linked by their overlapping ecological needs. Presuming that one of these types outcompetes the other, over a number of generations we will tend to see more of that type and less of the other. Eventually, we will likely see only the better-performing type in the population. This then looks just like natural selection, running until the fixation of a phenotype within the group. Yet when it is revealed that the two “phenotypes” are in fact different species, apparently this should change our assessment: we should say this is a case not of natural selection but of *competitive exclusion*. However, considering we have just shown that a group made of two species can meet the requirements for a paradigm Darwinian population as set out thus far, it may be appropriate to ask whether this distinction is very well motivated.

It is interesting to consider whether we could just accept the above case as an example of natural selection in at least some very abstract sense. However, it would seem to push things too far to claim it is a particularly paradigmatic case of natural selection, simply on the basis of its high α . After all, for all we have said, these two species could be extremely different from one another. The value of α can be high when there is only minimal overlap in the properties of the population members—just enough of an overlap such that “more of me means less of you,” which could result through competition over a single limiting resource. For example, in tidal areas or estuaries, bird species and aquatic species such as crabs or larger fish may compete over prey (see, e.g., Kneib 1982; Crowder, Squires, and Rice 1997). A “population” made of members of particular bird and crab species is simply not a group that will undergo paradigmatic evolution by natural selection as a whole, regardless of how strong the competition over food may be.

It is also important to note, however, that although there is something problematic with the idea that our gerrymandered group above is a paradigm Darwinian population, the fault does not lie with our analysis of α . The value of α is intended to be a measure of the interconnectedness of a population with respect to reproductive competition, and the population in this case exhibits a high level of this type of competition. Any member that has further offspring will have a strong negative effect on the reproductive success of the rest of the group, and this effect is distributed all but perfectly evenly. So this is a case in which the group deservedly has a high α . Therefore, the reason we should be wary of classifying this as a paradigm Darwinian population cannot be because it has a low level of reproductive competition. It looks as though some further criterion, not met in this case, is still required.

5. Addressing the Gap. Exactly why a group of two disparate species with high reproductive competition cannot be a paradigm Darwinian population raises subtle questions regarding natural selection. However, one clear issue

here involves how traits may spread through the population. If there are two markedly different species within a single putative population, either any new beneficial trait will only be able to spread through a portion of that group or an entire suite of different traits will have to spread with it.

This is because there are broadly two ways in which a fitness-enhancing trait might propagate through such a group over a series of generations. First, it might spread through the members of the species in which it arose. Even if the trait goes to fixation within that species, this will mean it only spreads through some portion of the population as a whole. Alternatively, the trait might give the organisms that possess it such an advantage over members of the other species, the species with the new trait displaces the other. In this case the trait might go to fixation within the gerrymandered population, but it will not be just this trait that has spread; the entire suite of traits of one species will have taken over the group. And as noted above, there may be very marked differences between our two species: in a “population” made up of particular species of birds and crabs, we will not see a trait ramify through the entire group so much as see the population go from a mixture of birds and crabs to a population of just birds. To put things less flippantly, this means that selection in a population that includes such radically different phenotypes will not lead to gradual, cumulative change, only wholesale change.

This leads us to a further concern. Since a population can have a high α even when its members are markedly different from one another, the respective niches of these subgroups may also differ in almost all properties apart from some important resource. In a case like this we will no longer see competition over a particular niche but only over that resource. To return to our example, if the crabs lose out to the bird species due to pressure on prey stocks, we will not see the crab niche taken over by a more fit organism, we will see it become (at least temporarily) empty.

It seems reasonable to require that for a group to constitute a Darwinian population, beneficial traits must be able to spread to fixation within the group. Furthermore, in paradigm cases of natural selection this should be able to occur without extreme alterations to either the entities within the population or the environmental pressures that drive selection on that population. We have just seen that this feature of paradigm Darwinian populations is not captured by considerations of reproductive competition alone, so something is missing from our account as it stands. In an attempt to rectify this gap, I will co-opt aspects of Alan Templeton’s notion of *exchangeability* (e.g., Templeton 1989, 1998).

5.1. Exchangeability. Templeton employs the idea of exchangeability as part of his “cohesion concept” of species. The cohesion concept states that a species is “the most inclusive population of individuals having the

potential for phenotypic cohesion through intrinsic cohesion mechanisms” (Templeton 1989, 12). In other words, a species is the largest population that has processes or structures in place that tend to make members of that population similar to one another. Note that the focus is on the mechanisms that produce phenotypic cohesion, rather than the cohesion itself. Templeton divides these mechanisms into *genetic* and *demographic* exchangeability.

Genetic exchangeability refers to the ability to combine genes with others in the group through sexual reproduction. This mechanism essentially characterizes the standard notion of “gene flow” and facilitates genetic relatedness and trait spread through the population. Demographic exchangeability refers to the extent to which organisms are matching “with respect to the factors that control and regulate population growth and other demographic attributes” (15). Templeton links demographic exchangeability with the notion of the organism’s *selective regime* (1998, 39). Here he deploys similar thinking to that behind Elliott Sober’s requirement of a “common causal influence”: that for natural selection to act on a population, the constituent individuals must be affected by the same selective forces (Sober 1984). Templeton defines this selective regime in terms of the extent to which the individuals share their *fundamental niche*. The fundamental niche of an organism delineates the ecological limits it can physiologically tolerate, such as the temperature and oxygen levels in the environment. This means organisms with high demographic exchangeability can be subject to the same types of environmental pressures as one another.

Genetic and demographic exchangeability are clearly different kinds of attributes, but they are importantly alike (at least for our purposes here) because they both promote a particular outcome, namely, within-group similarity. So exchangeability in this context means that the members of the population have some mechanism(s) that make(s) them likely to resemble one another in some important respects. Ability to interbreed will tend to lead to phenotypic similarity over successive generations, while demographic exchangeability means that population members are likely to be moulded by similar selective and developmental pressures. Templeton claims high exchangeability is crucial for natural selection to occur, because a drive toward cohesion allows new variants to spread throughout the group: “the cohesion mechanisms that define species status are therefore those that promote genetic relatedness and that determine the populational boundaries for the actions of microevolutionary forces” (Templeton 1989, 14). If a beneficial trait can only spread through part of the population due to a lack of genetic or demographic exchangeability, then this cannot be achieved.

I think it would be too quick to claim that this is the only way we might capture the missing element in our account of paradigm Darwinian populations, but there are a few reasons why this looks to be a suitable criterion to incorporate into the framework. First, although exchangeability was pre-

sented by Templeton as an account of species, the foundation of his account is based on evolutionary considerations. He reasons from the notion of an evolving lineage to the notion of species, rather than the other way around. Second, this is an account of species specifically intended to apply broadly across different kinds of organisms. As noted at the start of this article, we know the domain of organisms that can undergo paradigmatic natural selection may extend beyond species boundaries as narrowly construed. So broad applicability is a desirable feature for our purposes.

My proposal, therefore, is that we can employ exchangeability to supplement α in delineating paradigmatic Darwinian populations. That is, in order to form a population that can undergo paradigmatic natural selection, both high α and exchangeability are required.

5.2. Exchangeability as the Further Criterion. Recall that our further criterion must show why two markedly different species with overlapping resource requirements do not form a paradigm Darwinian population, without simply stipulating that all of the members of such populations must be conspecifics. First, it is crucial to note that invoking a criterion of exchangeability is not just the bare declaration that Darwinian populations can only contain members of the same species. The role exchangeability plays as the basis of Templeton's account of species is not important to us here. What is important is his link between exchangeability and natural selection. This means even if we accept that Templeton's criteria comprise the correct account of species, we will still have an explanation for why all paradigm Darwinian populations are made up of conspecifics. Namely, high exchangeability is required for paradigmatic natural selection, and as long as the cohesion concept of species is true, organisms with high exchangeability will necessarily be in the same species.

More importantly, if Templeton turns out to be wrong about what defines a species it would not matter for our account at all. Let us imagine Templeton is mistaken, and members of different species can have high exchangeability, or members of the same species can have low exchangeability. As far as our account is concerned, this would just mean that some populations made up of different species may undergo paradigmatic natural selection (as long as their exchangeability were sufficiently high) or that some populations made up of only one species may fail to be capable of undergoing paradigmatic natural selection (if their exchangeability were sufficiently low). So if high exchangeability turns out to be the correct account of species membership, our account at least tells us why conspecificity is associated with the formation of paradigm Darwinian populations, and regardless of this, we can adopt the criterion of exchangeability completely independent of its role in defining species.

Another nice feature of exchangeability as a further criterion is that it is largely dissociated from α . It is possible to have both groups with high

exchangeability but low reproductive competition (as in the case of isolated members of a single species) and groups with high reproductive competition but low (not zero) exchangeability (as in the case of two very different species that compete over a single crucial resource). Additionally, the addition of exchangeability to our account helps clarify the role interbreeding plays with respect to Darwinian populations. When present, genetic exchangeability is an important ingredient in delineating a paradigm Darwinian population, but if there is adequate demographic exchangeability present, it is not a necessary ingredient. So this is in keeping with Godfrey-Smith's claim that interbreeding is important, while exchangeability itself extends more broadly than this.

I now turn to the more central issue: how exchangeability deals with the earlier problem case. If we add exchangeability to our account, the example of a population that contains two different species under intense resource competition is no longer troubling and indeed produces some further positive results. Since the subgroups are stipulated to be composed of different species, it is likely that there is at least some difference in their fundamental niches, and in as much as this is the case, they will have nonmaximal exchangeability. To the extent that they are markedly different species, with markedly different fundamental niches and an inability to interbreed, the less our gerrymandered group will resemble a paradigm Darwinian population. Although α is high in such a case, exchangeability will be low.

What is more interesting, the higher the exchangeability, the more our criteria will classify this two-species group as a paradigm Darwinian population. And I think this is exactly what we would want. If exchangeability is high because of significant overlap of fundamental niches or the possibility of genetic mixing, and we know that there is intense reproductive competition, it seems we should not exclude this group merely on the basis of concerns regarding conspecificity.⁹

Here it may help to think of bacteria so we are not distracted by common assumptions regarding sexual reproduction. For example, in a paper by Hanage et al., the investigators differentiate a number of *Neisseria* species, in spite of the fact that all of the bacteria inhabit the human nose and exhibit lateral gene transfer. The authors are well aware that sorting organisms into different groups in spite of overlapping niche spaces and a shared gene pool is problematic, and they note that distinguishing species in such a setting "is a matter more of human interest and attention than any intrinsic evolutionary process" (Hanage, Fraser, and Spratt 2005, 6–7).

It might be that over time we will do away with such species divisions, but the important point is that as far as the possibility of natural selection is

9. Note that Templeton's account coincides with mine here, as according to his view, such a group simply ought to be considered a single species. Therefore, this type of case can only arise if we were to adopt an account of species at odds with Templeton's.

concerned, it seems any prior attention to species divisions will be at best a red herring. These strains of bacteria share genes and are under very similar selective regimes. In such circumstances, if current species distinctions state that the population contains different species, then so much the worse for the idea that a paradigm Darwinian population must all be of the same species.

To summarize: including exchangeability as part of our analysis of a paradigm Darwinian population means the case that appeared problematic for the account is dealt with in a way that categorizes it appropriately and has further supplementary benefits by classifying other cases more precisely. For perspective, I depict some pertinent cases in figure 3.

6. Alternative Accounts of Populations. At this point, I turn to consider two recent alternative accounts of biological populations. The first is a framework intended to give a quite general account of populations in biology, while the other argues for pluralism—the view that there is no privileged set of properties that pick out the populations that can undergo evolution. Note that although these positions appear to be at least contrary to one another, it is less clear that my position is in direct opposition to either. If there is a single account of biological populations to be had, it still may be useful to investigate the properties that delineate specific subtypes of these populations. Alternatively, if there is no uniquely privileged set of properties that pick out the

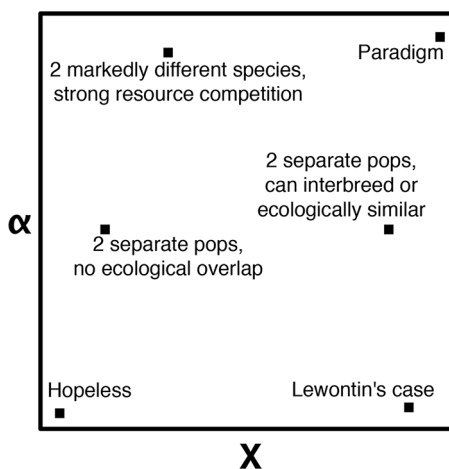


Figure 3. Cases according to strong competition (α) and exchangeability (X): previously problematic example of two different species under severe resource competition, two alternative realizations of population C from section 3, and Lewontin's case from section 2. I presume that exchangeability is reasonably high here since the bacterial colonies do fine in exactly the same broth.

biological populations, this does not rule out the possibility of saying something more specific regarding paradigm Darwinian populations in particular.¹⁰ In any case, comparisons between our views may help elucidate further issues. I begin with Roberta Millstein's account.

In two recent papers (2009, 2010) Millstein has set out "the causal interactionist population concept":

- Populations (in ecological and evolutionary contexts) consist of at least two conspecific organisms that, over the course of a generation, are actually engaged in survival or reproductive interactions, or both.
- The boundaries of the population are the largest grouping for which the rates of interaction are much higher within the grouping than outside. (Millstein 2010, 67)

In spite of our differing scopes, my account and Millstein's causal interactionist concept are clearly very similar. Most significantly, we both deny that proximity per se is important and instead treat interconnections between group members as of paramount significance. There are also some important differences, however.

For instance, we pick out different kinds of interconnections. While I take a specific kind of negative interaction as central and then include the presence of one or more cohesion mechanisms as a further criterion, Millstein groups survival and reproductive interactions together, where this subsumes reproductive competition as well as cooperative interactions. This seems apt, given that Millstein is concerned with ecology and evolution more broadly conceived, but how does it fit with my more narrow project?

I note two points here. First, cooperation seems to be neither necessary nor sufficient for a group to undergo paradigmatic natural selection. We know significant natural selection can occur without cooperation, and it is difficult to see how a population that only involved cooperative interactions could exemplify a paradigm Darwinian population, given my earlier remarks regarding the explanatory role of natural selection. Second, cooperation is certainly an important factor in certain instances of natural selection—by

10. This may also speak in part to Barker and Velasco's (2013) "deep conventionalism" regarding the population concept. At least as I read it, their primary concerns stem from the insight that there can be multiple legitimate claims regarding both what the population concept is meant to capture and also exactly what properties matter in delineating such populations. Since I restrict my target to just the populations that can undergo paradigmatic natural selection, restrict the relevant properties to reproductive competition and exchangeability, and (as becomes clear below) freely admit that there is no precise or objective cutoff to distinguish paradigmatic Darwinian populations from other Darwinian populations, I hope I avoid much of the force of Barker and Velasco's arguments.

increasing the fitness of particular types of individuals, for example—but this does not mean it should be included in our criteria regarding which groups are able to undergo paradigmatic natural selection. So it seems that an account specifically regarding populations that can undergo paradigmatic natural selection should include only competitive interactions, although cooperation will likely be an important feature when delineating populations for other purposes.

A further difference is that Millstein's account states that a population must be composed of conspecifics. I have outlined some reasons why we should resist restricting the analysis of paradigm Darwinian populations in this way, and here the difference between our accounts is perhaps significant. Given that Millstein is providing an analysis that extends more broadly than just populations that can undergo paradigmatic natural selection, if conspecificity is too restrictive for an account of these populations, it seems it will be too restrictive for her project also.

I now consider a very different approach, which directly opposes the project of giving a singular account of the populations that can undergo natural selection. In a recent article, Jacob Stegenga presents a number of arguments for what he calls a "radically pluralistic" population concept; an account that "denies both population uniqueness and population sharpness" (Stegenga 2014). For Stegenga, a lack of uniqueness means there are multiple sets of properties and weightings that could be employed to pick out the relevant groups, while a lack of sharpness means that even once a set of properties is chosen to delineate the populations, these properties will admit of degree and likely be vague at their boundaries.

Although Stegenga and I take substantially different approaches to this issue, I suspect we agree regarding many of the relevant points. Crucially, however, one of his claims is that an account that invokes causal connections will be unable to systematically pick out the populations capable of undergoing natural selection. If correct, this could substantially undercut the foregoing and so must be considered further.

Stegenga notes that if the relevant "causal connections" are specified too abstractly, the account will be too inclusive, allowing cases that should not be considered populations that can undergo natural selection. Alternatively, a more fine-grained specification of these connections will be very heterogeneous, and deciding which of these fine-grained causes we consider to be relevant will render the delineation of any population arbitrary. I hope to have shown here that this is not necessarily the case once we restrict our concern to paradigm Darwinian populations. The causal connection of reproductive competition is much narrower than the broad categories of "causally connected (somehow or other)" or even being "more fit," and in the course of this article we have progressively refined the view in order to ensure it is not too inclusive. However, reproductive competition is still quite a generally char-

acterized kind of causal connection, rather than a heterogeneous collection of fine-grained causal interactions such as “hiding better” or “growing higher.” In this way, we can capture a type of connection that finds a middle ground between being too inclusive and too fragmented: we have specified a particular type of connection in quite precise terms, while still allowing that this connection could be realized in multiple ways.

Unfortunately, the threat of arbitrariness still remains. Our framework for picking out paradigm Darwinian populations employs two properties, which raises questions regarding how to weigh these properties against one another. For example, we can imagine a situation in which drawing a population boundary one way picks out a group with a high level of reproductive competition but low exchangeability, while drawing the boundary another way picks out a group with a high level of exchangeability but low reproductive competition. Stegenga might be right that such a decision cannot be resolved in a nonarbitrary manner (although I think it might be premature to come to this conclusion just yet), so the best we can do in some cases might turn out to be a partial ordering over the candidate paradigm Darwinian populations. However, the situation is not as poorly behaved as the radical pluralist position appears to suggest. This is not a scenario of indefinitely many multidimensional constructs to choose from but a partial ordering generated through the use of two criteria. Even when it is not entirely clear which boundary will define the “most paradigmatic” Darwinian population in a given situation, at least we still know what we are looking for.

As noted above, Stegenga also argues that even if we were able to settle on a unified account of the types of causal connection that matter for picking out the relevant populations, specifying a boundary for some particular Darwinian population via its causal connections is potentially problematic. For example, will we only allow direct causal interactions or both direct and indirect interactions? The first seems too exclusive, likely to subdivide otherwise very tightly connected groups, while the latter seems too inclusive, incorporating groups such as the extreme case of ring species. Again, using the conceptual tool of directed, weighted graphs, I hope to have shown how a moderate view can negotiate this seeming dilemma. Indirect connections regarding reproductive competition are certainly relevant for population membership, but the more indirect or weak these connections, the less paradigmatic the Darwinian population will be.

Nevertheless, from this discussion we can see that like Stegenga’s position, my account also denies population sharpness: α and exchangeability are both graded quantities. A population does not exhibit reproductive competition or not; it exhibits reproductive competition to a certain degree, and the same observation holds regarding both genetic and demographic exchangeability. Given this graded nature of the phenomena, it would be artificial to specify any particular point where the account sharply declares a group of organisms to be

able to undergo paradigmatic natural selection. Rather, our analysis indicates that groups of organisms will form along a continuum of being paradigm Darwinian populations to a greater or lesser extent. This means that if we adopt α and exchangeability as the quantities that determine which groups can potentially undergo paradigmatic natural selection, again we may find this approach opens up the possibility for uncertainty regarding where the boundaries of these Darwinian populations “really” lie.

This is no failing of the framework, however. Populations will often lack strict boundaries and exhibit varying levels of interaction, and how powerfully natural selection acts on these populations will be similarly graded. In this case, we should not attempt to give a precise cutoff for the transition from mere grouping to Darwinian population to paradigm Darwinian population, as this would fail to reflect the biological facts.

The current account is compatible with pluralism regarding populations more generally, while still maintaining that in as much as we are concerned with paradigm Darwinian populations in particular, α and exchangeability will be the properties that define such populations. A lack of clear boundaries or complete ordering of candidate paradigm Darwinian populations does not mean we must reject the idea of a single account regarding the attributes exemplified by these populations, or the significance of those attributes with respect to the role of natural selection.

7. Summary. We have identified two quantities that determine whether a group of entities forms a paradigm Darwinian population. It is not enough to just stipulate that the members of such populations must all be in the same location and belong to the same species. What really matters is the presence of certain relations between the members of the group. One of these relations is the extent to which these individuals are in reproductive competition with one another. This can be represented by the variable α discussed by Peter Godfrey-Smith. After a few attempts, we settled on an analysis of this variable, employing the conceptual tool of directed, weighted graphs. However, an analysis of α is not enough to give us a full account of a population that is able to undergo paradigmatic natural selection. There is at least one further requirement: an adequate amount of exchangeability. I claim that if a group of organisms exhibits sufficient levels of α and exchangeability, then in the presence of phenotypic variation, fitness differences, and heredity, they will undergo paradigmatic natural selection.

I then compared my account with two recent views regarding biological populations in the literature. Both of these accounts are concerned with the possibility of a more general characterization of populations and appear to be in direct opposition to one another. The framework I propose here is compatible with either, and I believe it also advances the discussion further: it allows us to say something informative about a particular, important type

of biological population, by adopting a grain of analysis that is neither too broad nor too narrow.

We might ask whether there are other requirements that a group of entities must meet before they form a group that has the potential to undergo paradigmatic natural selection. I have not said anything that would rule this out, but neither have we encountered any examples that would suggest α and exchangeability do not give us a good account of such populations. So far, these two criteria appear to classify cases correctly, and we have uncovered a number of interesting issues on the way.

REFERENCES

- Barker, M., and J. Velasco. 2013. "Deep Conventionalism about Evolutionary Groups." *Philosophy of Science* 80 (5): 971–82.
- Crowder, L., D. Squires, and J. Rice. 1997. "Nonadditive Effects of Terrestrial and Aquatic Predators on Juvenile Estuarine Fish." *Ecology* 78 (6): 1796–1804.
- Godfrey-Smith, P. 2009. *Darwinian Populations and Natural Selection*. New York: Oxford University Press.
- Hanage, W., C. Fraser, and B. Spratt. 2005. "Fuzzy Species among Recombinogenic Bacteria." *BMC Biology* 3 (6). doi:10.1186/1741-7007-3-6.
- Kneib, R. T. 1982. "The Effects of Predation by Wading Birds (Ardeidae) and Blue Crabs (*Callinectes Sapidus*) on the Population Size Structure of the Common Mummichog, *Fundulus Heteroclitus*." *Estuarine, Coastal and Shelf Science* 14:159–65.
- Lennox, J., and B. Wilson. 1994. "Natural Selection and the Struggle for Existence." *Studies in History and Philosophy of Science* 25:65–80.
- Lewontin, R. 1970. "The Units of Selection." *Annual Review of Ecology and Systematics* 1:1–18.
- Millstein, R. 2006. "Natural Selection as a Population-Level Causal Process." *British Journal for the Philosophy of Science* 57:627–53.
- . 2009. "Populations as Individuals." *Biological Theory* 4:267–73.
- . 2010. "The Concepts of Population and Metapopulation in Evolutionary Biology and Ecology." In *Evolution since Darwin: The First 150 Years*, ed. M. Bell, D. Futuyama, W. Eanes, and J. Levinton, 61–86. Sunderland, MA: Sinauer.
- Nanay, B. 2005. "Can Cumulative Selection Explain Adaptation?" *Philosophy of Science* 72 (Proceedings): 1099–1112.
- . 2010. "Natural Selection and the Limitations of Environmental Resources." *Studies in History and Philosophy of Biological and Biomedical Sciences* 41:418–19.
- Neander, K. 1995. "Pruning the Tree of Life." *British Journal for the Philosophy of Science* 46:59–80.
- O'Malley, M., and J. Dupre. 2007. "Size Doesn't Matter: Towards a More Inclusive Philosophy of Biology." *Biology and Philosophy* 22:155–91.
- Sober, E. 1984. *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Chicago: University of Chicago Press.
- . 1995. "Natural Selection and Distributive Explanation." *British Journal for the Philosophy of Science* 46:384–97.
- Stegenga, J. 2014. "Population Pluralism and Natural Selection." *British Journal for the Philosophy of Science*. doi:10.1093/bjps/axu003.
- Templeton, A. 1989. "The Meaning of Species and Speciation: A Genetic Perspective." In *Speciation and Its Consequences*, ed. D. Endler and J. Otte, 3–27. Sunderland, MA: Sinauer.
- . 1998. "Species and Speciation: Geography, Population Structure, Ecology, and Gene Trees." In *Endless Forms: Species and Speciation*, ed. D. J. Howard and S. H. Berlocher, 32–43. New York: Oxford University Press.