

Beanbag Holobionts

by

Caleb Hazelwood

Department of Biology
Duke University

Defense Date: April 1, 2025

Approved:

Dan McShea, Advisor

V. Louise Roth

Alex Rosenberg

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in the Department of Biology in
The Graduate School of Duke University
2025

ABSTRACT

Beanbag Holobionts

by

Caleb Hazelwood

Department of Biology
Duke University

Defense Date: April 1, 2025

Approved:

Dan McShea, Advisor

V. Louise Roth

Alex Rosenberg

An abstract of a thesis submitted in partial fulfillment of the requirements for the
degree of Master of Science in the Department of Biology in
The Graduate School of Duke University
2025

Copyright by
Caleb Hazelwood
2025

Abstract

In this paper, I consider two arguments concerning the status of holobionts as evolutionary individuals—one *rejects* the thesis by privileging the “stability of lineages” (*sensu* Godfrey-Smith 2009) and the other *supports* the thesis by privileging the “stability of traits” (*sensu* Veigl et al. 2022). I argue that the tension between these two arguments arises from two fundamentally different accounts of natural selection. I suggest that each account of selection corresponds to a unique account of evolutionary individuality. This strategy entitles us to a modest pluralism: holobionts are evolutionary individuals on one account of selection but not on the other.

Contents

Abstract.....	iv
1. Introduction	1
2. The Squibrio.....	3
3. Stability of Lineages.....	6
4. Stability of Traits	9
5. Dueling Darwinisms.....	12
6. The Two Nature of Selection	17
7. Conclusion: In Defense of Beanbag Holobionts	23
References.....	25

1. Introduction

What is a biological individual? There are several good ways of answering this question (e.g., Clarke 2010; Pradeu 2016). In this paper, I am specifically interested in what it means to be an *evolutionary* individual. A common starting point for understanding evolutionary individuality is natural selection.

Evolutionary individuals are often defined as those entities that are “visible” to natural selection—the entities upon which natural selection can act (Clarke 2010). They are sometimes called the “units of selection” (e.g., Gould and Lloyd 1999).

It seems there is unanimous agreement that evolutionary individuals should be carved out by natural selection.¹ The consensus crumbles, however, when we try to determine which entities fit the bill. Philosophers of biology are no strangers to this debate: Williams’ argument against group selection (1966) and Dawkins’ argument against anything but genic selection (1976) are familiar cases. But as the horizons of biology expand, so do the candidates for the units of selection. In recent years, philosophers have argued that natural selection may

¹ This is not to say that natural selection is the only relevant evolutionary process. That is, delineating evolutionary individuals with natural selection does not entail adaptationism. Evolutionary individuality, as I understand it, only demands the *capacity* to evolve via selection. A population of individuals capable of evolving under selection can also drift, mutate, migrate, and exhibit any other evolutionary process you like.

not only act on genes, organisms, groups, and species, but on “holobionts,” as well.

A holobiont is a multispecies consortium—the sum of a plant or animal host and the entire community of microbes that lives in and on it. Holobionts are of particular interest to philosophers of biology because, as we learn more about how our microbiomes are inherited within and across generations, they put pressure on our notion of what it means to be an evolutionary individual. This paper asks and attempts to answer the question: are holobionts evolutionary individuals? My answer will be yes and no—it all depends on the account of “natural selection” one has in mind. I will arrive at this conclusion with the help of a case study, which I introduce in the following section.

2. The Squibrio

In 2005, Robert Brandon wrote the following passage:

Every day moths get eaten by birds. Now I'm as fond of moths as the next person, but generally this fact about moths does not upset me. However, in an otherwise insightful article published almost 20 years ago Beatty (1984) unleashed into the philosophical literature a population of moths that confused him and that seem to have caused considerable confusion since. It is time to clarify this confusion and to let these poor moths rest in peace. (2005, p. 162)

Another twenty years later, and Beatty's moths are finally in peaceful repose. But a new population of model organisms from a thought experiment gone awry has invaded the literature. I refer to a population of squid—the Hawaiian bobtail squid, *Euprymna scolopes*, to be exact—and its symbiotic partners, the bacteria *Vibrio fischeri*.

The “squibrio” (“squid” + “*Vibrio*”), as this symbiotic association has come to be called, is the darling of the holobiont literature. It is an especially interesting case because it exhibits characteristic features of evolutionary individuality in some ways and, in other ways, is far too ephemeral to be “seen” by selection.

Some holobiont researchers are tempted to consider the squibrio an evolutionary individual because it appears to bear a holobiont-level adaptation: bioluminescence. The body cavity of the squid contains a “light organ,” an internal chamber that has evolved to be inoculated exclusively by *Vibrio fischeri* (Visick and McFall-Ngai 2000). When the bacteria aggregate in large numbers within the squid, they are bioluminescent, lighting the squid aglow. This is adaptive for the nocturnal squid, as predators in the water column below mistake the squid for moonlight. It is also thought to be adaptive for the bacteria, given that free-living colonies (i.e., those that do not reside in the light organs of squid) lose their bioluminescence (Miyashiro and Ruby 2012). Thus, it is the emergence of a holobiont-level trait—bioluminescence—that tempts some to attribute the status of evolutionary individuality to the squibrio.

But others are tempted to say that the squibrio is *not* an evolutionary individual. This is because the squid-bacteria associations are ephemeral; the relationship between the host’s lineage and the lineages of its symbionts only lasts a single generation. As with many holobionts, the bacteria are transmitted horizontally, not vertically. That is, a squid’s symbionts are not inherited from the squid’s parents but are obtained from a communal stock to which many squid—even unrelated neighbors—contribute (Visick and McFall-Ngai 2000). Therefore, the squibrio does not form tidy parent-offspring lineages. It is this lack

of a unified lineage that tempts other holobiont researchers to deny the status of evolutionary individuality to the squibrio.

Just like Beatty's moths, the squibrio has caused considerable confusion about what it means to be an evolutionary individual. A surfeit of literature has deployed this case study to argue for or against the claim that holobionts are units of selection. Often—and to the dismay of microbiologist readers—these arguments paint the biological details in broad brushstrokes, reducing the squibrio to rhetorical caricature.

I propose that enough ink has been spilled in vain over these poor squid. In this paper, I will argue that philosophers of biology have not been able to agree over the squibrio's status as an evolutionary individual for one crucial reason: they are relying on two different accounts of natural selection—one that privileges the stability of lineages, and another that privileges the stability of traits. With different accounts of selection in hand, they arrive at entirely different notions of evolutionary individuality. I will consider each of these accounts in the following sections, then advocate for a modest conceptual pluralism that acknowledges the practical utility of each.

3. Stability of Lineages

The first view I will consider is the view that evolutionary individuals must form parent-offspring lineages. Call this view the “Stability of Lineages” (Veigl et al. 2022). Its most comprehensive defense is from Peter Godfrey-Smith (2009). Godfrey-Smith’s lineage-centered account of evolutionary individuality builds on Richard Lewontin’s formalization of natural selection—a tripartite recipe with three ingredients: variation, inheritance, and differential reproduction (1970). As Godfrey-Smith notes, an important upshot of Lewontin’s recipe is that, in abstracting away from contingent biological facts, his formalization allows for the possibility of natural selection at levels above and below the individual organism. Any entity that possesses these three features is a unit of selection, an evolutionary individual.

Godfrey-Smith defines a “Darwinian population” as “a collection of causally connected individual things in which there is variation in character, which leads to differences in reproductive output (differences in how much or how quickly individuals reproduce), and which is inherited to some extent. Inheritance is understood as similarity between parent and offspring, due to the causal role of the parents” (2009, p. 40). The “causally connected individual things” that exhibit the features to satisfy Lewontin’s recipe may, therefore, be called “Darwinian individuals” (ibid., Godfrey-Smith 2013).

Darwinian individuality is Godfrey-Smith's answer to the question of evolutionary individuality. On his account, reproduction is the name of the game. Reproduction of Darwinian individuals from other Darwinian individuals is necessary for a population to possess the features of inheritance and fitness. Simply put, Darwinian individuals are the kinds of things that make family trees. Parent-offspring lineages are necessary because they are the basis of evolutionary fitness. To say that I am fitter than my neighbor is to say that I am likely to have more offspring than her; to say this, we must be able to—at least in principle—delineate which offspring are hers and which are mine.

Therefore, when we privilege the Stability of Lineages, it follows that many holobionts are not evolutionary individuals. This is because, as mentioned earlier, many holobionts acquire their symbionts via horizontal transmission—from sources other than their parents. Consider the squibrio:

... if you are a squid, there is no mechanism ensuring that the bacteria in you are the offspring of bacteria in your parents, or any other specific individuals. The bacteria in you might come from many sources, and some might have not been inside squid for many generations. Squid-Vibrio combinations “make more of themselves” in *one* sense, but not in the sense that gives rise to parent-offspring lineages. The parent-offspring lines connect only the parts – they connect bacteria with bacteria and squid with

squid... So the combinations are not Darwinian individuals... They are a metabolic knotting of reproductive lineages that remain distinct. (2013)

This is sometimes referred to as the “too many parents” objection: “... there must be some parent–offspring similarity, and the clarity of a ‘parent-offspring’ relation of the relevant kind is inversely related to the number of parents—if there are too many parents there are no parents at all.” (Godfrey-Smith 2012; see also Doolittle and Inkpen 2018). Since the squibrio commits this violation, it is not a Darwinian individual, and therefore not an evolutionary individual.

4. Stability of Traits

Other researchers are unsatisfied with the solution that the Stability of Lineages provides: accounts such as Godfrey-Smith's deny many holobionts—especially those with microbiomes that are facultative rather than obligate and are transmitted horizontally rather than vertically—the status of evolutionary individuality.

Some such dissenters of Godfrey-Smith's account have proposed an alternative: rather than privilege the stability of *lineages* in demarcating evolutionary individuals, we should instead privilege the stability of *traits*. Among them, Veigl, Suarez, and Stencel have given the most recent defense (Veigl et al., 2022) of the claim that establishing a lineage is a *sufficient but unnecessary* condition for establishing heredity.

Veigl et al. “contend that contemporary evolutionary approaches to heredity are severely limited because they fail to explain some cases of what [they] call Stability of Traits.” They define the Stability of Traits as a “phenomenon manifested by the transgenerational persistence of variation in the phenotypic characteristics of some individuals in a population at a specific level of the biological hierarchy, not necessarily forming a continuous line of ancestry” (ibid.). They acknowledge the Stability of Traits as a “functional concept” that captures the inheritance of phenotypic variants—not only those that are

transmitted through parent-offspring lineages, but also those that are *reconstituted* by parts from myriad lineages.

These authors propose the “reconstitutor” as an alternative to Godfrey-Smith’s “Darwinian individual.” They define the reconstitutor as the “structure resulting from a set of relationships between different elements or processes that are actively involved in the recreation of a specific phenotypic variant in each generation regardless of ... whether they stand in a continuous line of ancestry” (ibid.). The reconstitutor, they argue, is the basic unit of heredity. It can be realized in many ways—e.g., a trait may be reconstituted via parent-offspring lineages, but this is not necessary. What matters most is that the trait of interest reappears from one generation to the next.

We see this kind of reconstitution in many holobionts. In fact, Veigl et al. call holobionts “exemplars” of the reconstitutor concept. Citing Godfrey-Smith, they take issue with identifying the “Darwinian individual as the unit that captures what needs to be preserved to detect hereditary relations,” and bemoan that “several researchers have rejected the idea that there are hereditary relations among holobionts” (ibid.). This rejection, they write, derives “from the narrow view of the process of inheritance encouraged by the assumptions of [the Stability of Lineages]” (ibid.). The authors offer their account of the reconstitutor as a solution to the squibrio imbroglio:

Imagine that, for some reason, a jointly expressed trait, like bioluminescence, shows transgenerational phenotypic variation. In this case, there is a phenomenon of [Stability of Traits] at the system level. But since the system does not constitute the type of structure that Godfrey-Smith considers necessary to acknowledge the existence of a unit of heredity, he would say that this phenotypic variation needs to be carved up either at the level of the squid or at the level of the *Vibrio*, as they are the parts of the consortium forming lineages. This is, however, problematic, for this would precisely mask the actual level where hereditary relationships are manifesting. (Veigl et al. 2022)

It's true, the squibrio may fail to demonstrate one or two necessary ingredients for Lewontin's recipe, and this absence may disqualify it from achieving the status of "Darwinian individual." But since the holobiont-level trait—bioluminescence—is capable of varying, and since those variations are capable of being reconstituted in successive generations of squibrio symbioses, it follows that the squibrio is a reconstitutor, and therefore an evolutionary individual.

5. Dueling Darwinisms

My own view is that the reconstitutor—Veigl et al.’s proposed “evolutionary unit of heredity”—is really no unit of heredity at all. Ultimately, a coherent account of heredity demands that its units are situated in parent-offspring lineages. To successfully argue for this claim demands more space than I can afford.

For present purposes, I am not concerned with the content of either position. I am, instead, concerned with a higher-order question: What’s gone wrong? How have we found ourselves in a situation where biologists and biologically oriented philosophers can disagree so fundamentally about whether selection acts on certain entities? In this section, I suggest an answer to this question: the disagreements over the boundaries of an evolutionary individual arise not in virtue of different assumptions about concepts like “reproduction” and “heredity.” The rift runs much deeper than that. Instead, disagreements over the boundaries of an evolutionary individual arise due to differing assumptions about *natural selection itself*.

Recall that we have defined an evolutionary individual as any entity upon which natural selection can act. Thus, the empirical task of individuating individuals transmutes into the philosophical task of defining natural selection. As we have established, Lewontin’s tripartite “recipe” for selection (and its

conceptual scion, Godfrey-Smith's account of "Darwinian individuality") has become the canonical definition: natural selection is heritable variation in reproductive fitness.

But one person's modus ponens is another's modus tollens. Say I endorse the following inference:

1. If it doesn't evolve by natural selection, then it's not an evolutionary individual.
2. (Most) holobionts don't evolve by natural selection.
3. Therefore, (most) holobionts are not evolutionary individuals.

This is the kind of inference we find in arguments for "Darwinian individuality." Because holobionts seldom demonstrate heritable variation in reproductive fitness (at the level of the entire holobiont, not its individual parts), they do not evolve by natural selection. Since they do not evolve by natural selection, they are not evolutionary (i.e., Darwinian) individuals. But this sort of inference can be countered.

Let us say, instead, that I find the conclusion of the previous inference unpalatable. Let us say that I am committed to the idea that holobionts are evolutionary individuals (for the reasons considered in Section 4). In that case, rather than affirming the antecedent, I deny the consequent.

1. If it doesn't evolve by natural selection, then it's not an evolutionary individual.
2. (Most) holobionts *are* evolutionary individuals.
3. Therefore, (most) holobionts *do* evolve by natural selection.

When one takes such a tack, they arrive at the conclusion that holobionts evolve by natural selection. They arrive at such a conclusion because they have agreed with the first premise—it would deny a biologist's credibility to assert otherwise—and they have taken the second premise on board as an assumption.¹ The person who prefers this second inference, the modus tollens, must therefore find a way to make it so that holobionts evolve by natural selection. But when one assumes Lewontin's tripartite recipe for natural selection, defending the conclusion of the modus tollens becomes quite tricky. The strategy for the proponents of the modus tollens has thus been to target the *components* of the recipe, e.g., by taking issue with lineage-centric notions of "heredity" and accounts of "reproduction" that demand material overlap (e.g., Charbonneau 2014; Papale 2021).

¹ Again, this is not to deny that populations of evolutionary individuals can evolve via drift, mutation, etc. But they must also have the *capacity* to evolve under selection.

This is the wrong way to go about defending the modus tollens. It forces us into making claims such as “heredity does not require parent-offspring lineages.” Such claims are frankly not viable. Parent-offspring lineages are a necessary ingredient in the recipe of selection; otherwise, there is no such thing as one lineage being “fitter than” another.

So, if that is the wrong way to defend the modus tollens, then is there a right way? There is at least a better way. The way to defend the modus tollens is to argue that there are legitimate ways of understanding natural selection beyond the tripartite recipe. That is, the proponent of the modus tollens ought to argue that there is not a biconditional relationship between “natural selection” and the accounts offered by Lewontin and Godfrey-Smith. Our revised first premise would look something like this:

$$\sim(\text{NS}_{\text{recipe}} \vee \text{NS}?) \rightarrow \sim\text{EI}$$

where “NS” means “natural selection” and “EI” means “evolutionary individual” (or, if you like, “unit of selection”). In informal terms: if an entity does not evolve by natural selection—where natural selection is understood to refer to *either* Lewontin’s tripartite recipe *or* some other legitimate account(s)—then that entity is not an evolutionary individual. Now, with such a conditional, one can argue a plausible modus tollens: holobionts *are* evolutionary individuals,

so holobionts must be subject to natural selection, bearing in mind that “subject to natural selection” is a disjunctive property that we have yet to flesh out.

The question now is whether we can find an alternative conception of natural selection—one that allows us to salvage the claim that holobionts, despite their lack of parent-offspring lineages, are subject to such a process.

6. The Two Nature of Selection

Is there an alternative conception of natural selection on offer? There is, and we may find it by looking to a decades-old debate in the philosophy of biology: the debate between the “causalists” and the “statisticalists.” For the sake of space, I must ride roughshod over many important and interesting details of the debate. (For a recent summary of and contribution to the literature, see Pence 2021). For our purposes, the important distinction is that statisticalists argue for a second conception of natural selection—one that arose with the advent of the “modern synthesis,” a historical period that saw the application of statistical tools to evolutionary theory and resulted in the construction of population genetics models.

According to the statisticalist, natural selection is a statistical summary of the changes in frequency of a trait of interest at the population level. It says nothing about the success of an organism’s lineage, an organism’s fitness, etc. According to the population genetics models of the modern synthesis, the statisticalist argues, natural selection *just is* the directional change in the frequency of a trait, as predicted by that trait’s fitness.

In one of the more recent arguments for statisticalism, Denis Walsh and colleagues distinguish the two senses of selection in the following way:

“The process of selection that Darwin postulated is in essence the change in lineage structure that occurs when there is variation in [individual] fitness. Accordingly, we shall call this phenomenon ‘Darwinian selection’ (or ‘D-selection’).”

“The other process, in which populations change in their trait distribution as a function of variation in their trait fitnesses, we’ll call ‘modern synthesis selection’ (or ‘MS-selection’)...” (Walsh et al. 2017).

Note that, by defining D-selection in terms of variation in the fitnesses of individual lineages, Walsh et al. are aligning D-selection with the tripartite recipe. In other words, D-selection is what Lewontin and Godfrey-Smith mean when they use the term “selection;” it is the process that picks out Godfrey-Smith’s “Darwinian individuals.” MS-selection, on the other hand, is a less stringent definition. It does not require individual lineages, reproduction, or inheritance. All it requires is that traits change in relative frequency as a function of their fitness over evolutionary time.

My ameliorative proposal is that, by introducing the distinction between D-selection and MS-selection, both camps of the holobiont debate—those that privilege the Stability of Traits *and* the Stability of Lineages—find a friendly conceptual framework. In other words, I propose that we distinguish between

two senses of evolutionary individuality, with each sense grounded in one of the two interpretations of “natural selection” just discussed.

Following Walsh et al. (2017) and having identified both D-selection and MS-selection as extant interpretations of “natural selection,” we should recognize corresponding notions of evolutionary individuality for each interpretation: D-individuals and MS-individuals.

D-individuals are the entities that evolve by D-selection. They satisfy Lewontin’s tripartite recipe because they possess heritable variations in reproductive fitness. They reproduce to create more D-individuals, thereby forming parent-offspring lineages. Conveniently, D-individuality—an abbreviation of “Darwinian individuality”—aligns perfectly with Godfrey-Smith’s account of the same name. Godfrey-Smith prefigures this analysis, in fact, when he writes that:

...the idea of a Darwinian population is not set up in a way that is specific to some particular trait. Breaking organisms into traits is treated as something that comes later in an analysis. What is recognized first is the collection of individuals with all their features, both repeatable and unique. (2009, p. 40)

MS-individuals, on the other hand, are the entities that evolve by MS-selection. They do not (necessarily) satisfy Lewontin’s tripartite recipe because they need not reproduce nor inherit traits from previous generations. MS-

individuals are merely the bearers of a trait of interest. How the trait arises is irrelevant so long as the trait is reliably reconstituted in each generation.

This dual notion of individuality—each corresponding to an account of natural selection grounded in the biological literature—lets us have our cake and eat it too. We can agree with Godfrey-Smith that most holobionts are not individuals *in the Darwinian sense* because they do not form parent-offspring lineages. They *may be* individuals, however, in a sense that is relevant for the modern synthetic interpretation of natural selection: they bear traits of interest.

At this point, I must waylay the objection that I am foisting an account of MS-individuality onto those who prioritize the Stability of Traits when, in reality, they hold no such view. That is correct: the critics of Darwinian individuality I cite have never explicitly assented to the “statisticalist” picture of natural selection nor its corresponding account of individuality. In fact, I suspect that many of them would balk at the idea. But that is why my proposal is an ameliorative one. My claim is that *if* those who prioritize the Stability of Traits continue to play by the rules set by those who prioritize the Stability of Lineages—e.g., (Veigl et al. 2021) and (Godfrey-Smith 2009), respectively—*then* they will lose the game every time. The “recipe” account of natural selection forces its critics to take positions such as, “heredity does not require parent-offspring lineages, only the reconstitution of traits.” But such a view is

incoherent: if heredity is the product of inheritance, then a trait must be inherited *from* one individual *by* another (or, more specifically, reconstituted with parts *from* one holobiont *by* another holobiont). Is this “from-one-by-one” relationship not the formation of a parent-offspring lineage?

I am not making the descriptive claim that those who prioritize the Stability of Traits consciously hold such a view. Instead, I am proposing that they ought to hold such a view. They require an account of natural selection—and a corresponding account of individuality—that backgrounds the individual lineages and, instead, privileges the frequency of the traits of interest. I am proposing that such an account is necessary to make sense of the “Stability of Traits.”

Some interlocutors have suggested that my argument collapses into a *reductio ad absurdum* against the Stability of Traits. The *reductio* assumes that statisticalism is false. The argument runs something like this: “The only way to make sense of the Stability of Traits is through the lens of statisticalism. But statisticalism is false. Therefore, there is no viable way to make sense of the Stability of Traits.” But this *reductio* fails. It fails because *if* statisticalism is false (which is my view), *then* it is false because it gets the metaphysics of natural selection wrong. The metaphysics of natural selection—and, more specifically, its status as a “genuine causal process”—are what is at stake in the debate between

statisticalists and causalists. No one denies the statisticalists' characterization of "MS-Selection," i.e., no one denies that, at least *in practice*, selection is often identified as the directional change of trait frequencies within a population, nor do they deny that population biologists predict evolutionary change via "trait fitness," a population-level statistical parameter. And it is this aspect of statisticalism—its characterization of MS-selection and not its metaphysics—that I am deploying here. In other words, MS-selection (as I have introduced it) is silent on metaphysical claims such as "natural selection is [or is not] a genuine causal process." Therefore, the *reductio* would have to be modified: "The only way to make sense of the Stability of Traits is through the lens of MS-selection—the trends captured by a set of modeling techniques upon which statisticalists build their argument for the claim that natural selection is not a genuine causal process. But statisticalism is false. Therefore, there is no viable way to make sense of the Stability of Traits." Clearly, the conclusion of our revised *reductio* does not follow. There is no reason why MS-selection cannot be understood as a causal process (which, again, is my view).

7. Conclusion: In Defense of Beanbag Holobionts

This shift in explanatory focus (from lineages to traits) calls to mind the quibble between Mayr and Haldane on the usefulness of “beanbag genetics” (Haldane 1864). Mayr’s criticism was that the practice of beanbag genetics—i.e., mathematically modeling *allelic* evolution at one locus by likening “the genetic contents of a population to a bag full of colored beans” (Mayr 1963)—eclipsed the importance of *organismal* processes (such as development and physiology) and grossly oversimplified genetic expression. Treating populations as bags of beans of different frequencies—reducing populations of organisms to simple statistical, allelic models—proved useful for empirical tractability but dismal for explanatory depth. Or so Mayr argued.

But now we know better. Models and other explanatory tools are constructed to be adequate for certain purposes (Parker 2020). Information about individual development, physiology, and the complexities of gene expression are necessary for explaining certain phenomena. But if what I want is a population-level apparatus that enables me to predict future frequencies of a trait of interest, then beanbags will do the trick.

An analogous argument may be run for evolutionary individuality. One account of evolutionary individuality need not be better than another in any holistic or absolute sense. If what I wish to understand is how competing *lineages*

have ramified in response to a selection pressure, then D-selection is necessary to tell the story. If, on the other hand, I wish to understand how competing *traits* have changed over evolutionary time, and I can abstract away from the details of individual lineages that underlie the shifting balance of traits, then MS-selection is the relevant model.

Veigl et al.'s "reconstitutors" are beans in a bag. What matters for their model is the trait, and the frequency with which the trait reappears in future generations. The underlying causal story of how that trait comes to be does not matter. As Haldane defended beanbag genetics, we may similarly defend beanbag holobionts.

References

- Brandon, R. N. (2005). The difference between selection and drift: A reply to Millstein. *Biology and Philosophy*, 20, 153-170.
- Charbonneau, M. (2014). Populations without reproduction. *Philosophy of Science*, 81(5), 727-740.
- Clarke, E. (2010). The problem of biological individuality. *Biological theory*, 5(4), 312-325.
- Dawkins, R. (1976). *The selfish gene*. Oxford university press.
- Doolittle, W. F., & Inkpen, S. A. (2018). Processes and patterns of interaction as units of selection: An introduction to ITSNTS thinking. *Proceedings of the National Academy of Sciences*, 115(16), 4006-4014.
- Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. Oxford University Press, USA.
- Godfrey-Smith, P. (2013). Darwinian individuals.
- Godfrey-Smith, P. (2012). Darwinism and cultural change. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1599), 2160-2170.
- Gould, S. J., & Lloyd, E. A. (1999). Individuality and adaptation across levels of selection: how shall we name and generalize the unit of Darwinism?. *Proceedings of the National Academy of Sciences*, 96(21), 11904-11909.
- Haldane, J. B. S. (1964). A defense of beanbag genetics. *Perspectives in Biology and Medicine*, 7(3), 343-360.
- Lewontin, R. C. (1970). The units of selection. *Annual review of ecology and systematics*, 1(1), 1-18.
- Mayr, E. (1963). *Animal species and evolution*. Harvard University Press.
- Miyashiro, T., & Ruby, E. G. (2012). Shedding light on bioluminescence regulation in *Vibrio fischeri*. *Molecular microbiology*, 84(5), 795-806.
- Parker, W. S. (2020). Model evaluation: An adequacy-for-purpose view. *Philosophy of Science*, 87(3), 457-477.

- Papale, F. (2021). Evolution by means of natural selection without reproduction: revamping Lewontin's account. *Synthese*, 198(11), 10429-10455.
- Pence, C. H. (2021). *The causal structure of natural selection*. Cambridge University Press.
- Pradeu, T. (2016). The many faces of biological individuality. *Biology & Philosophy*, 31, 761-773.
- Veigl, S. J., Suárez, J., & Stencel, A. (2022). Rethinking hereditary relations: the reconstitutor as the evolutionary unit of heredity. *Synthese*, 200(5), 367.
- Visick, K. L., & McFall-Ngai, M. J. (2000). An exclusive contract: specificity in the *Vibrio fischeri*-*Euprymna scolopes* partnership. *Journal of Bacteriology*, 182(7), 1779-1787.
- Walsh, D. M., Ariew, A., & Matthen, M. (2017). Four pillars of statisticalism.
- Williams, G. C. (1966). *Adaptation and natural selection: A critique of some current evolutionary thought* (Vol. 1). Princeton university press.