

The Selectionist Rationale for Evolutionary Progress

Hugh Desmond

Forthcoming in Biology and Philosophy

Abstract

The dominant view today on evolutionary progress is that it has been thoroughly debunked. Two important theoretical reasons are seen to support this view: (1) natural selection provides no rationale for progress, and (2) natural selection need not even be invoked to explain large-scale evolutionary trends. In this paper, I challenge both reasons, and by analysing how natural selection acts in heterogeneous environments, propose that natural selection over macroevolutionary scales provides a rationale for a pattern of evolutionary unfolding, characterised by radiation across an increased range of exploitation of environmental heterogeneity.

KEYWORDS: Evolutionary Progress – Natural Selection – Laws – Evolutionary Contingency – Environmental Heterogeneity

1. Introduction

In the not too distant past, claims that life had made some type of ‘progress’ throughout evolutionary history were considered nearly self-evident, even by prominent biologists:

During the past billion years, animals as a whole evolved upward in body size, feeding and defensive techniques, brain and behavioral complexity, social organisation, and precision of environmental control. (...) Let us not pretend to deny in our philosophy what we know in our hearts to be true (Wilson 1992, p. 187).

Even at the time when Wilson made that statement, evolutionary progress was viewed either with “disdain” (Sober 1994, p. 19), or at least as thoroughly debunked. Today, this stance remains fundamentally the same. While sweeping views of life’s history receive considerable attention via the literatures on evolutionary trends (McShea 1994; McShea and Brandon 2010) and on major transitions (Calcott and Sterelny 2011; Maynard Smith and Szathmary 1995), the concept of evolutionary progress itself no longer receives much attention from philosophers of biology or philosophical biologists. The main academic areas where the concept of

evolutionary progress finds an intellectual home are either among anthropologists interested in understanding why people spontaneously believe in ‘higher’ and ‘lower’ organisms (e.g. Atran 1998), or among science educators interested in weaning students off similar beliefs (e.g. Johnson et al. 2012; Werth 2012).

One important class of reason why ‘progress’ remains a no-go is social-ethical in nature, for instance, how concepts of progress were used to justify the old eugenics (see Desmond forthcoming). Another class of reason is empirical: several of the most important concepts of evolutionary progress have not received much empirical support (McShea 1998, 2001a). However, this paper will focus on the theoretical underpinnings of progress, and in particular on the argument that progress *lacks* any such underpinnings. One reason lies in the nature of natural selection itself. Echoing earlier statements by G. C. Williams¹, S. J. Gould writes:

The bare-bones mechanics of the theory of natural selection provides no *rationale for progress* because the theory speaks only of adaptation to local changing environments (Gould 2002, pp. 467–468, my emphasis).

An example of such a ‘rationale’ would be the following: evolutionary history has ‘progressed’ towards a greater ability to process environmental information because it allows an organism to respond more adaptively to an environmental change and hence is selectively favoured across many if not all environments. So what Gould and Williams are saying here is not that such a selectionist rationale is false per se – it may be true that some kind of information processing ability is consistently favoured across selective environments – but rather that the causal structure of natural selection itself provides no reason to believe that it is true. I call this the *groundedness problem*: no selectionist rationale for progress seems to be grounded in the nature of natural selection.

Besides the groundedness problem, detractors of evolutionary progress often strengthen their case by pointing to what could be called the *redundancy problem*, where selection is claimed to be simply superfluous for explaining evolutionary trends, since the trends for which there is empirical support (like the increase of maximum body size or complexity) can be explained as resulting from a species’ evolution resembling a random walk, or a ‘diffusion from the left wall’:

Since space remains available away from the left wall and toward the direction of greater complexity, new species occasionally wander into this previously unoccupied domain, giving the bell curve of complexity for all species a right skew, with capacity

¹ “There is nothing in the basic structure of the theory of natural selection that would suggest the idea of any kind of cumulative progress.” (Williams 1966/2018: 35)

for increased skewing through time (...) Such an extension of the right tail will occur in a regime of entirely random motion for each item, so long as the system begins at a wall (Gould 1996, pp. 105–106).

I will later show how this ‘neutral’ view of trends is distinctively mathematical (Huneman 2010; Lange 2013), and because of this is often treated as the null for explaining passive trends. This leads to the redundancy problem, which holds that selectionist explanations are superfluous because they are insufficient to dislodge this null explanation.

Taken together, the groundedness and redundancy problems seem to undermine selectionist rationales for large-scale evolutionary trends. By extension, selectionist rationales for evolutionary progress are undermined: not only are such rationales not grounded in natural selection itself (and any claim that they are just evinces a misunderstanding of natural selection), but they are explanatorily useless. Hence, any argument that progress has resulted from selection automatically becomes suspect (i.e., what really motivates a biologist to invoke selection if a random walk can do the trick?).

The goal of this paper is to argue that there is at least one selectionist rationale for evolutionary progress that is (1) grounded in the causal structure of natural selection, and (2) not explanatorily redundant. The core proposal is that environmental heterogeneity – including the “local changing environments” in the citation above – does not simply act as a source of contingency, but rather gives rise to selective processes that causally explain a large-scale trend towards species that are better able to exploit environmental heterogeneity. This then re-opens the question to what extent value-neutral concepts of progress are viable from a fundamental theoretical perspectives. Such concepts construe ‘better’ and ‘worse’ in selectionist terms such as more and less competitive, or superiorly and inferiorly adapted. I will discuss two concepts that seem viable: unfolding progress and competitive progress.

There are two caveats about the scope of the argument. First, some may reject value-neutral progress concepts for *a priori* reasons, and hold that the word “progress” essentially involves the imposition of societal and/or ethical values. Addressing this issue is beyond the scope of this paper; suffice it to say that the value-neutral usage is not uncommon (e.g. Conway Morris 2003; Dawkins 1997; Dennett 1995; Mayr 1994), and it was this usage that Gould aimed to undermine in the quote above. The question for this paper is whether there is any grounded selectionist rationale for it. The second caveat is this paper will not take a stand on whether or not evolution actually “progresses”. The focus is on the *coherence* of rationales for progress with fundamental theoretical principles – not on the truth of such rationales. Claims that evolution has progressed may be wrong for a variety of reasons, both empirical and *a priori*;

the modest argument of this paper is only that, if they are wrong, it is not because the selectionist rationales for them cannot cohere with the causal structure of natural selection.

The paper is structured as follows. In the following second section I explain what is intended by a ‘selectionist rationale for progress’. In the third and fourth sections I analyse and argue against the neutral view of evolutionary trends. In the fifth section I distinguish between two processes of selection in heterogeneous environments, only one of which can be said to be a source of evolutionary contingency on the largest scales. In the sixth section I make the case for a selectionist rationale for a trend in the exploitation of environmental heterogeneity. This overcomes the groundedness problem, and in closing I suggest what a selectionist rationale for progress could look like.

2. The Structure of Selectionist Rationales for Progress

Claims about evolutionary progress are vast generalisations about the history of life and the patterns instantiated throughout that history. Often such claims have also concerned how life could have *possibly* evolved (Rosslénbroich 2006; Shanahan 2001). Since the danger exists that such claims fall short of all sorts of epistemic standards – and devolve into mere speculation – it is important to analyse the structure of these claims.

2.1 Laws of Evolutionary History. In claims that evolutionary history has progressed, or progress-claims for short, only a very particular aspect of that history is of interest, whether complexity, body size, genetic information, and so on. Hence a measure M is introduced that assigns a real number to any given lineage in evolutionary history (Figures 1a, b). In this way a time-slice of evolutionary history can be represented as a frequency distribution, which changes over time (Figure 1c). A *pattern* P is a further abstraction of this frequency distribution, obtained by isolating one particular statistical property of this distribution, and mapping its changes over time (Figure 1d). For instance, a pattern can describe how the maximum of the frequency distribution changes over time. Patterns can also concern either the minimum or the average – or any statistical property for that matter (skew, kurtosis, etc.). When the pattern describes a series of increases in the statistical property, it is called a *trend*.

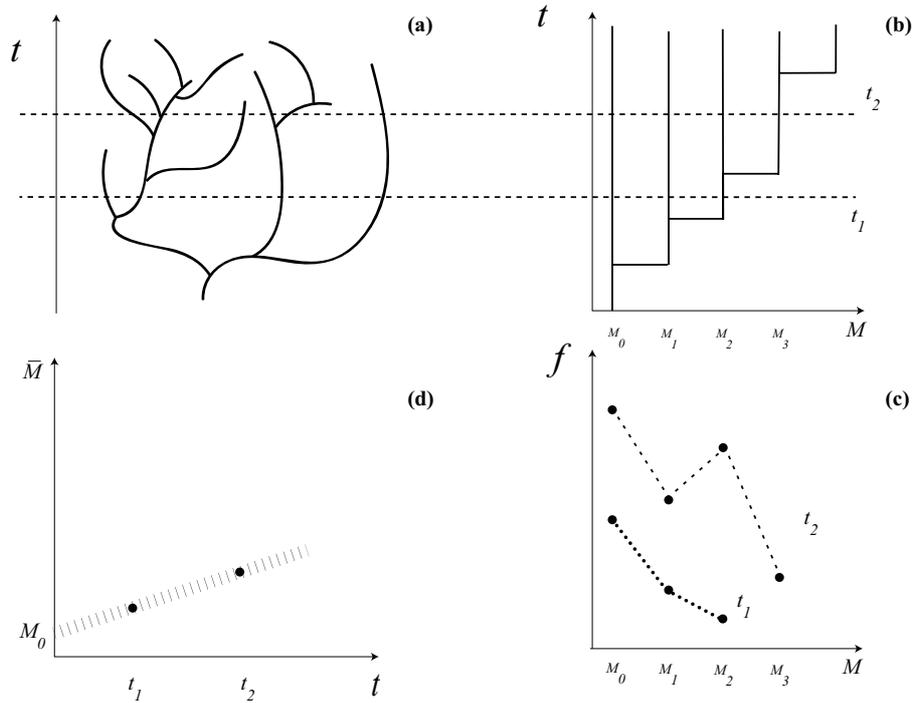


Figure 1: Patterns in Evolutionary History (clockwise from top left): (a) an evolutionary history; (b) the occupation of M -states over time; (c) absolute frequency distributions of lineages over M -states at times t_1 and t_2 ; (d) an increase in average M across all extant lineages over time.

The type of progress-claim of interest in this paper will involve a trend of a certain counterfactual robustness, a *law of evolutionary history* with the following structure:

“All possible evolutionary histories instantiate a pattern P.”

Here, an ‘evolutionary history’ is the sum total of lineages or ancestor-descendant series, both extant and extinct. Since large-scale studies in palaeontology concern generalisations at the level of species (e.g. Payne et al. 2009), an evolutionary history can also be represented as the history of species, where individual organisms standing in ancestor-descendant relations to each other have been grouped together according to species.

Implicit here is that the space of possible evolutionary histories is restricted to *biologically* possible histories. The initial conditions of life may be varied (different environments, different concentrations of compounds) but the laws of physics are the same, the periodic table is still the same and so on. The space of evolutionary histories more or less corresponds to “replays of life’s tape”, even though the formulation is more general, since (potential) evolutionary histories on other planets are included as well.

Even when restricted to biologically possible histories, the dominant view on laws of evolutionary history is that they do not exist. Most generalisations are deemed at most ‘rules’, like Cope’s Rule which states that populations in the same lineage tend to increase in body size

over time. Other generalisations may be exceptionless, like Dollo's law (a lineage that has evolved will never return to its initial state; Dollo 1893) or the Zero Force Evolutionary Law (McShea and Brandon 2010), but these generalisations are consequences of statistics (as well as structural factors such as the size of the space of biological possibilities) rather than the dynamics of biological processes such as natural selection (see McShea and Brandon 2010, p. 95).

This thought is encapsulated by the Evolutionary Contingency Thesis, or ECT (Beatty 1995):

All generalisations about the living world: (a) are just mathematical, physical, or chemical generalisations (...) or (b) are distinctively biological, in which case they describe contingent outcomes of evolution (Beatty 1995, pp. 46–47).

Applied to laws of evolutionary history, the ECT implies that any pattern P that is instantiated across all possible evolutionary histories (i.e., non-contingent patterns) must be the consequence of theories in mathematics, physics, or chemistry. There is no pattern that is the consequence of distinctively biological processes, such as genetic mutation, drift, or natural selection. The ECT implies that there cannot be a selectionist rationale for a law of evolutionary history, as this would be a distinctively biological generalisation describing a non-contingent outcome.

2.2 Selectionist Rationales. Let us now define selectionist rationales with more precision. A 'rationale' is not a prediction of what patterns occur in actual evolutionary history. As I will understand it – and this, I believe, is compatible with Gould's usage as well – a rationale denotes a causal process that is grounded in an underlying theoretical principle. In particular, a 'selectionist rationale for X, entailed by Y' denotes how Y grounds a selective process that leads to X. Thus, a selectionist rationale is a relation between a causal term (Y) and an outcome term (X).

This allows a more precise formulation of the groundedness problem, which states: natural selection cannot ground a selective process that leads to evolutionary progress. And in fact, most historical selectionist rationales for progress are not grounded in natural selection itself, but in speculative claims that certain types of traits (body size, information processing capacity, etc.) are selectively favoured across most if not all environments. Beatty's discussion of the ECT echoes the groundedness problem, since he explicitly lists changes in environment

as one of the sources of evolutionary contingency (Beatty 1995, p. 53): this is why the ECT implies there is no selectionist rationale for progress.²

While a rationale can be counteracted in empirical evolutionary history, it does predict definite patterns under suitably idealised conditions. The idealised pattern associated with a selectionist rationale can be called a *selectionist law of evolutionary history*:

Given effectively random environments and without long-term constraints on natural selection, all evolutionary histories instantiate a trend in a measure M that is caused by selection for increased M .

Here evolutionary histories are idealised so that any causal process arising from the ‘bare-bones mechanics of selection’ is translated into an empirical trend. Note that this implies that in actual evolutionary history, which may not have suitably idealised conditions, there may not be an observable trend in M even though there is a selectionist rationale for increased M .

Before proceeding, it is important to unpack the idealised conditions. The first is that evolutionary histories are constrained to those with effectively random environments. This means that any resulting trend caused by selection is a consequence of selection’s causal structure, rather than a consequence of particular environments that happened to occur frequently. For Gould (as well as Beatty), this condition actually disqualifies any expectation of a selection-caused trend, since the (mistaken) assumption is that random succession in environments means random selective processes.³

The second idealised condition – no long-term constraints on natural selection – is intended to ensure that any genetic or phenotypic constraints in populations would be overcome over time scales of tens to hundreds of millions of years. Thus, even if, for instance, entirely different *Baupläne* survived the Cambrian–Ordovician extinction event (see Gould 1989), a selectionist rationale for M would eventually cause a trend in M to occur. While this specification of no-constraints is common enough (the special formulation of the Zero Force Evolutionary Law (ZFEL) makes use of it: see McShea and Brandon 2010, p. 3), and while it underlies the position of strong adaptationism (see Godfrey-Smith 2001), it should be noted

² For this reason, there being a (grounded) selectionist rationale would seem to contradict the ECT in spirit, even though the ECT could still be saved in letter since selectionist rationales cannot be used for exceptionless generalisations over actual, empirical evolutionary history.

³ See e.g. “The sequence of local environments in any one place should be effectively random through geological time - the seas come in and the seas go out, the weather gets colder, then hotter, etc. If organisms are tracking local environments by natural selection, then their evolutionary history should be effectively random as well.” (Gould 1996, pp. 139–140)

that some short-term constraints are necessary for there to be any historicity at all, since otherwise eusocial organisms could suddenly appear in a prokaryotic world.⁴

As a side remark, note that a ‘selection-caused trend’ is not the same as a ‘selection-driven trend’, which is the common term in the literature on largest-scale evolutionary trends (following McShea 1994). The qualifiers ‘driven’ and ‘passive’ strictly denote statistical properties rather than causes. In a driven trend, increases in the measure M are more frequent than decreases in a lineage; in a passive trend, increases and decreases occur with equal frequency. Not all driven trends are caused by selection: some may be caused by developmental ratchets (e.g. Saunders and Ho 1976), to the accumulation of components (e.g., the ‘neutral ratchets’ in Lukeš et al. 2011). Not all driven trends have specific causes: the (special formulation of the) ZFEL results from the statistical claim that it is more probable that diversity increases than decreases. Conversely, a selection-caused trend can be both driven and passive. (The neutral view, discussed in sections 3 and 4, assumes that all passive trends can be explained without selective causes.)

2.3 Selectionist Rationales for Progress. How does progress figure in this picture of selectionist rationales for evolutionary trends? A relatively common view is that the concept ‘evolutionary progress’ is strictly non-scientific, involving the grafting of values onto a scientific concept (i.e., directionality). This stance view is encapsulated by Sober’s pithy formula: “progress = directional change + values” (Sober 1994, p. 20). The view also receives expression in Gould’s colourful verdict that progress is “a noxious, culturally embedded, untestable, nonoperational, intractable idea” (Gould 1988, p. 319).

If this view were obviously true, then all progress-claims could be dismissed without further ado (within the contexts of biology and philosophy of biology). However, it is not obviously true – and this may also be why Gould himself puts so much effort into undermining the theoretical ground for progress-claims. Upon closer inspection, the link between natural selection and progress is much more intricate than would seem from Sober’s formula alone.

This link is not apparent if natural selection is simply conceptualised in terms of the Lewontin conditions⁵ – heritable fitness variation. Rather, the link between selection and progress only can make sense when one adopts a concept of natural selection as involving ‘competition’ of some sort, and/or potential to produce adaptation. Arguing for this would

⁴ How precisely constraints on natural selection should be understood is quite a vexed issue, beyond the scope of this paper. I thank <BLINDED> for pressing me on this point.

⁵ As, by the way, Sober explicitly does in (Sober 1984).

require a separate paper; suffice to say that, in a group of orca and krill, there is heritable phenotypic variation and associated fitness variation, but no natural selection. Note also that Darwin himself primarily wrote about the “struggle for life”, which is quite different from heritable fitness variation (for extensive argument, see Lewens 2010), and that some recent accounts of natural selection conceive of it as involving reproductive competition (Godfrey-Smith 2009, pp. 49–53).

This is important, because when selection is an adaptation-producing and/or competitive process, the outcomes of a selection-caused trend can be thought of as the ‘victors’ and/or the ‘superiorly adapted’ relative to a given environment. Evolution by natural selection produces ‘better’ outcomes in the following value-neutral sense: they are defined relative to the ‘preferences’ of natural selection, not relative to human moral or societal preferences. The relevant parameters could be better competition, or better adaptation.

This linking of natural selection and progress was present in many historical accounts of evolutionary progress, which relied on an argument that some property – whether complexity, cognition, or something else – allowed general adaptive or competitive superiority across environments (for an overview, see Rosslénbroich 2006). It is present in more recent notions of competitive progress’ (e.g. Vermeij 1987, pp. 419–421) or ‘adaptive progress’ (e.g. Dawkins 1997; Mayr 1994) These specific concepts of progress were perhaps ungrounded – the victors and superiorly adapted are only ‘better’ relative to an environment, and environments change⁶ – but it does show how the value-neutral usage of ‘progress’ is legitimate and not to be dismissed by default.

This intimate link between selection and progress has been off the table for quite some time, largely because the groundedness and redundancy problems have been seen as insurmountable. If selection cannot ground progress anyway, then Sober’s formula may seem accurate. However, overcoming these problems revives the value-neutral usage of the term ‘progress’, and opens up the possibility that evolutionary trends could also be considered progressive. This is the reason why this paper is not merely about selectionist rationales for evolutionary trends, but also about selectionist rationales for evolutionary progress.

To sum up this section, there are three levels of claim that should be distinguished: progress-claims about evolutionary history; selectionist rationales for those claims; and the theoretical grounding for selectionist rationales. The groundedness problem states that there is

⁶ So for these particular versions of (selective) competitive and adaptive progress I agree, with Gould and Williams, that natural selection offers no rationale for them. If a progress-claim is to be viable, it must involve different concepts of progress.

no theoretical grounding for selectionist rationales; the redundancy problem states that, even if there would be theoretical grounding, selectionist rationales are explanatorily superfluous anyway. The following two sections tackle the redundancy problem; sections 5 and 6 address the groundedness problem.

3. The Neutral View of Evolutionary Trends

This section sketches the received view on the redundancy problem: passive trends can be explained parsimoniously as neutral trends. Before a general characterisation, consider the trend in nestedness as an example (McShea 2001a, 2001b). The measure of nestedness, or “hierarchical object complexity” (McShea 1996), has four values – prokaryote, eukaryote, multicellular, colony.⁷ Evolutionary history started at the prokaryote level, and maximal nestedness increased monotonically, reaching the upper limit some 800 million years ago (Figure 2). It is not conclusively known whether this trend is driven or passive (see McShea 2001b): for the following, we will assume it is passive.

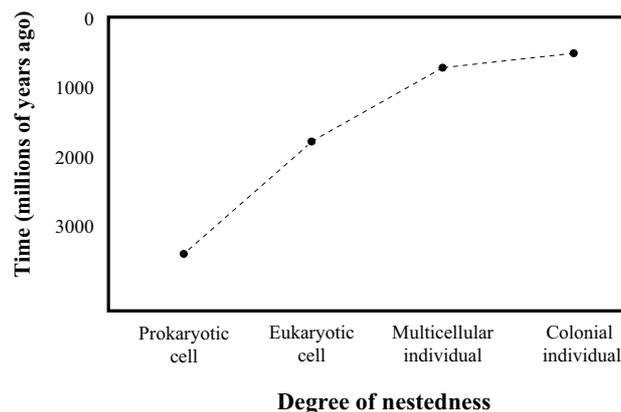


Figure 2. The increase in degree of nestedness across evolutionary history. (Redrawn from McShea 2001b.)

Is selection needed to explain this increase in maximum nestedness? According to the neutral view, it is not: since life began at the minimum value of nestedness, from that starting point only more nested life could evolve, even if this evolution were a random walk. All one needs is for species to “occasionally wander into previously unoccupied domain” (Gould 1996, pp. 105–106); that, together with starting at the minimum value, is all one needs to generate a passive trend in nestedness.

This parsimonious explanation generalises for any scalar measure M . A random walk through M -space, where states are defined as values of M , can be represented as a game where

⁷ This taxonomy can be further fine-grained according to the internal differentiation and interaction between parts (see McShea 2001a). This need not further concern us here.

the player flips a coin and then moves a checker left if the coin lands heads, or right if it lands tails.



Figure 3: Random walk through M -space (adapted from Sober 1994).

Even with a fair coin it is very probable that a player will have a run of tails, or a sequence where tails far outnumber heads. Hence, if this game is played by a large number of players simultaneously, the expectation is that the checkers will on average move to the right over time.

In an evolutionary context, the coin flip is an idealisation of the sum total of processes affecting the lineage evolution over a certain unit of time. The time step for M -increase (T_{\uparrow}) defines a time-scale where the probability of an M -increase is $\frac{1}{2}$ (and similarly for the time step for M -decrease, T_{\downarrow}). A coin flip involves only a single time step ($T_{\uparrow} = T_{\downarrow}$) and is simply the length of time it takes for a person to pick up a coin and flip it. The time-step for changes in nestedness is on the scale of tens to hundreds of millions of years when M ; for changes in body-size, the relevant time-scale is millions of years.

This explanation of passive trends in M thus relies on nothing but statistics, i.e., the evolution of the binomial distribution as one conducts an increasing number of Bernoulli experiments. It is a distinctively mathematical explanation (Huneman 2010; Lange 2013) of a biological phenomenon, and hence, not surprisingly, it generalises for any passive trend in measure M . This is why Gould can apply it not just to increases in complexity, but also to directional changes in body size or information processing capability: all occur for the same reason as happens if enough people are flipping coins for long enough, the checkers will move from the initial state to the right. This vindicates Beatty's ECT: largest-scale passive trends are just not interesting evolutionary phenomena, and do not warrant any biological (let alone selectionist) explanation.

Since this is a sweeping explanation of (passive) evolutionary trends, I call this the *neutral view* of trends, in analogy with the neutral theory of molecular evolution (Kimura 1968) and the neutral theories of biodiversity (Hubbell 2001). It is inspired by, but should be distinguished from, the quantitative 'MBL model', constructed by Tom Schopf, S.J. Gould, David Raup, Dan Simberloff, and others in the 1970s. The MBL model represents evolutionary changes in lineages as occurring at fixed intervals with preset probabilities, regardless of how well individuals in that lineage were adapted to their environment (for an in-depth discussion, see Sepkoski 2016 or Turner 2011). In one version of the model, extinction and speciation

events were modelled (with the goal of explaining changes in diversity: Raup et al. 1973); in another version, changes in morphology were modelled (Raup and Gould 1974). The neutral view takes the basic modelling assumptions of MBL models, applies them to all passive evolutionary trends, and claims that this is the ‘null’ explanation of passive trends.

In sum, rather than a narrowly defined quantitative model, the neutral view is a sweeping explanation of all passive trends that moreover positions itself as the null view: not only is it more parsimonious than selectionist explanations, but it also avoids speculation about what selective environments did and did not occur (such that it also avoids the groundedness problem). Given that there are no documented largest-scale driven trends, this deprives selectionist rationales of playing any necessary function in explaining largest-scale evolutionary trends.

However, upon a closer look, the neutral view depends on an assumption that is speculative in the same way as assumptions underlying problematic selectionist rationales for progress.

4. Problems with the Neutral View

In this section we will focus more closely on the conditions that justify the neutral view of passive trends. Of particular interest are the two following assumptions: (1) that $T_{\uparrow} = T_{\downarrow}$ (so that the evolutionary coin is fair, so to speak), and (2) that T_{\uparrow} is not too large. The first condition is a restatement of the explicit assumption that lineages evolve as random walks. The second condition is necessary, because if, for instance, T_{\uparrow} were to be on the scale of hundreds of billions of years, the neutral view would not predict any observable trend given that life is three billion years old. We will now consider in detail what these two conditions mean when translated in terms of underlying microevolutionary processes (genetic mutation, phenotypic plasticity, selection, drift).

To get a grasp on these processes, let Ω_g be the space of all possible genotypes in which the evolution of lineages can be represented.⁸ Limiting the length of possible genomes to some large number N , the size of Ω_g is unknown but finite, and by all accounts, many times larger than all the protons in the universe (Wright 1932). Other commentators have simply denoted the size of Ω_g as “vast” (Dennett 1995) or “hyper-astronomical” (Kauffman 1993).

⁸ Phenotypic change can occur without genetic change (due to phenotypic plasticity), so this concept of a space of possible genotypes is just a heuristic. See critical discussion of biological possibility in (Maclaurin and Sterelny 2008).

Another way of thinking about the vastness of Ω_g is that if genotypes were sampled randomly, only a vanishingly small subspace of Ω_g would ever be explored in evolutionary time⁹. So if we assume that M -states supervene on Ω_g -states¹⁰ – partitioning Ω_g – then given the vastness of Ω_g and each M_i region, a lineage randomly sampling genotypes in the M_0 region would most likely never reach the M_I region. Even the empirically observed time-step of 10-100 million years for increases in nestedness (McShea 2001a) is much shorter than what one would expect if individual lineages were to randomly sample Ω_g states.

Thus, for all practical purposes, it is basically impossible that a random walk through M -space could be realised by random walks in the lower level space Ω_g (or more precisely, the probability is exceedingly low). So, for the neutral view to be plausible, lineage evolution in the vast lower level space must be biased. The sources of bias are well documented. One important type results from corridors of functional equivalence dictated by the structure of genotype-phenotype maps (even as sampling within those narrow corridors is random: Kimura 1968). Another type is mutational biases (as discussed in Louis 2016). A final type of bias that is indispensable for explaining the evolution of phenotypes, is that of natural selection spreading adaptive traits.

Note thus that for the neutral view to work, it needs selection operating at the microevolutionary level. Nor is this an inconsistency: what matters is that selection at the microevolutionary level does not translate into a bias in the evolution through M -space. Thus, the preceding discussion allows us to pinpoint the following crucial assumption underlying the neutral view:

In geological time-scales, where the succession of environments is effectively random, M -favouring and M -disfavouring selective environments occur with equal (but not too small) frequencies for any M .

⁹ Louis 2016 explores the case of protein folding: if a protein were to sample random spatial configurations before finding the right one, it would take in the order of 10^{120} times the age of the universe. (The number of possible protein configurations is about 10^{150} , proteins can sample about 10 trillion (10^{13}) configurations per second, and the age of the universe is in the order of 10^{17} seconds.) The number of possible genotypes is vastly larger than 10^{150} : for haploid genotypes consisting of 1000 genes and 2 possible alleles per gene, there are $2^{1000} \cong 10^{316}$ possible genotypes. Even this is an underestimation of Ω : the number of coding genes is variable, can be much larger than 1000 (humans have around 20.000 of them by latest estimations), many organisms are diploid, many genes polymorphic, and so on. In any case, even if evolution were to consist of a billion lineages, the sampling rate of genotypes is much slower than 10 trillion times per second. Thus, random sampling of genotypes would correspond to the exploration of a vanishingly small subregion of Ω .

¹⁰ This is likewise a simplifying assumption. For instance, wolfpacks, or cooperating predatory bacteria (*Myxococcus*), are considered to exemplify multicellularity (Berleman and Kirby 2009; Grosberg and Strathmann 2007). Thus, changes in level of nestedness may be environmentally induced and underdetermined by genotype alone.

Note the similarity with Gould's remarks about environments in geological time-scales in footnote 3.

What is problematic about this assumption is that it needs to hold for any M . Recall why it is necessary: only in this way can the neutral view represent passive trends in any M as resulting from random walks in M . However, if one analyses more closely what 'effectively random environments' and 'any M ' mean – that is, what reference classes are involved – then it becomes clear that one cannot evaluate the truth or falsity of this assumption. In other words, it is speculative, and thus an unsuitable basis for the sweeping conclusions implied by the neutral view.

First let us give more precision to what a 'random' environment means. One could say that an environment is random_1 when it is picked randomly from the set of all metaphysically possible environments. This is an exceptionally large set, and much larger than would be appropriate in the context of a discussion of evolutionary history, because it includes environments where magic and Lewisian tiny miracles happen. More appropriate would be understanding an environment to be random_2 when it is picked randomly from the set of "biologically possible" environments, which are the environments that could be expected to occur during a replay of life's tape (see section 2). Such biologically possible environments share all the following with actual environments: the laws of physics, the periodic table, all other regularities in chemistry. What may be different is that asteroids may impact the Earth at different times or not at all; there may be different patterns of global temperature change; the compositions of the atmosphere and the lithosphere may be different; different developmental constraints (like body plans) may come to dominate early on in evolution; and so on.

By contrast to possible environments, when it comes to possible M , the neutral view does not allow for any limitations. The distinctively mathematical structure of the neutral view implies that the measure M can be drawn from a much larger space of possibility than simply the biologically possible M . The neutral view thus also applies to exaggeratedly artificial M , such as "the size of horse-sized butterfly-shaped wings", or "the number of fluorescent feathers that can self-propagate". Are these traits, on the whole, selectively neutral across biologically possible environments? There are obvious developmental constraints that prevent evolutionary history from realising any increases in such measures, but there seems to be no good reason why we should not think that such massive and weird wings may be selectively *disfavoured* in most biologically possible environments. One can think of many metaphysically possible environments, replete with magic and miracles, where weird wings may be selectively favoured. However, this is not the case for biologically possible environments: such wings may

be selectively disfavoured in most biologically possible environments. In other words, the neutral view's assumption that all measures are (effectively) selectively neutral cannot be interpreted as an unproblematic, fall-back assumption. It has the same speculative status as assuming that some measures are (effectively) selectively favoured.

This point holds for more biologically relevant measures, like body size. Some have argued that body size is selectively favoured more often than disfavoured (Bonner 1988), but others have put forth equally good arguments showing that increased body size can put an organism at a selective disadvantage (e.g. McShea 1998). What is the exact proportion of (biologically possible) environments that favour increased body size in a given lineage versus the environments that disfavour size increases? This is a question that no amount of palaeoclimatological, palaeogeological, or palaeoecological investigation will help answer, because such investigation at best only reveals what the proportions of environments were in *actual* evolutionary history. Nobody knows, and we should refrain from assuming not only that-favouring environments occur more frequently, but also that they occur with equal frequency. Claiming body size is selectively neutral is as speculative – and in this sense, ungrounded in natural selection itself – as claiming it is consistently selectively favoured.

In conclusion, the neutral view seems to offer an elegant, distinctively mathematical explanation of passive trends in any measure M , thus showing any selectionist rationale for the trend to be entirely superfluous. However, for the neutral view to work, all M must be effectively selectively neutral, but there is no ground for either believing or disbelieving this assumption. In this way, the neutral view is guilty of the same speculation about selection in possible replays of life's tape that some selectionist rationales are guilty of. Or in other words, the neutral view ultimately faces the same groundedness problem as the selectionist rationales it aims to undercut: it makes speculative assumptions about what selective environments occur.

5. Selection in Heterogeneous Environments

Let us now turn to the groundedness problem. The proposal here is to revisit the claim that environmental heterogeneity is a source of contingency for evolution by natural selection (Beatty 1995). The underlying assumption that selection occurs within well-defined, constant environments is an idealisation that is not only always false, but more importantly, often not justified. This section shows how, once the assumption is relaxed, various selective dynamics in response to EH can be identified. The next section argued that these entail a selectionist rationale (for a trend in EH-exploitation) that avoids the groundedness problem. The final section will connect this with progress.

5.1 Local and Extended Environments. What does “local” precisely refer to when saying that natural selection concerns only adaptation to “local changing environments” (Gould 2002, pp. 467–468)? Elsewhere Gould goes no further than describing it as the “immediately surrounding environment” (Gould 2011/1996: 139), and illustrates it by noting how the woolly mammoth evolved a coat as an adaptation to a cold environment, but is not in any way a “cosmically better or generally superior elephant” (Gould 2011/1996: 139).

Yet, as is well known, changes in the immediately surrounding environment can also occasion specific selection pressures. For instance, when the magnitude of fluctuations exceeds the tolerance of individual phenotypes, then environment tracking (e.g., phenotypic plasticity) or mixed strategies (e.g. polymorphism) are selectively favoured. Which one of the latter two is favoured then depends on the spatiotemporal scale of fluctuation: if the temporal scale exceeds generation time or the spatial scale exceeds the typical habitat range of an individual, then mixed strategies are favoured.¹¹

Such adaptations to environmental heterogeneity (EH) are not adaptive to EH as such, but rather to specific *patterns* of EH (not to be confused with a pattern of evolutionary history). Thus, avian migratory behaviour is adaptive to seasonal variation, but not to variation in weather (time-scale of hours or days), nor to climatic variation (time-scale of tens to thousands of years). Induced defences of bryozoans are adaptive only when the density of predators varies on a time-scale that is longer than the reaction time, but shorter than the lifespan. The motility of bacteria is adaptive only when nutritional resources vary on a spatial scale that is equal to or smaller than the habitat dimensions.

This means that the “local” environment is not literally the immediately surrounding environment – it is, by virtue of the selective pattern of EH, extended across space and/or time. Figure 5 illustrates this. Heterogeneity here is slower fluctuation in some variable X , which could represent temperature, concentration of nutrients, presence of predators, and so on. Environment E_1 is characterised by a slow fluctuation in X (large period T_X); environment E_2 is characterised by a fast fluctuation in X (small period T_X).

¹¹ For the details of the selective models underlying these conclusions, see Levins 1968, or discussion in AUTHOR forthcoming.

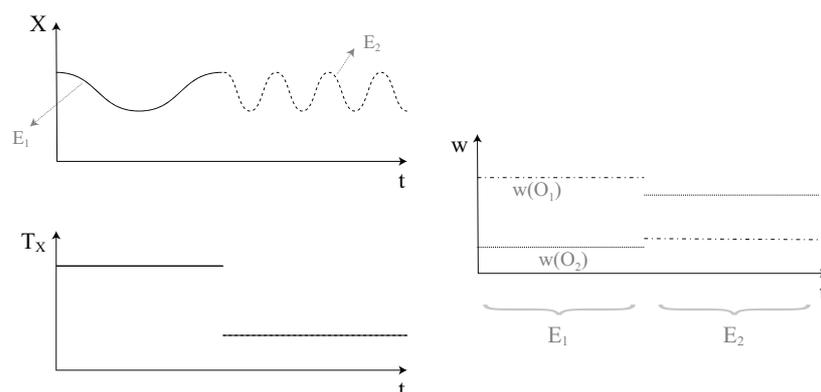


Figure 4: Patterns of EH can individuate selective environments. The two figures to the left are different representations of two environments, one with a long-period variation in X (E_1), and another with a short-period variation in X (E_2). The figure to the right shows how differences in period lead to difference in fitness in two organism types.

In this way, the selective environments that favour adaptations to heterogeneity are *necessarily* extended over time (and/or space), since the pattern that occasions the selection pressure is extended throughout time (and/or space).

5.2 Adaptations to Novel Environmental Heterogeneity. When selective environments are defined by patterns of EH they not only become necessarily extended in space and/or time, but the boundaries of the selective environment also become ambiguous. What precisely defines the selective environment: must it be defined in terms of the patterns an organism has already encountered, or can it involve novel patterns the organism has never encountered before? This ambiguity points to a distinction between two types of adaptation to environmental heterogeneity. An example will help introduce a more general characterisation.

Consider migratory behaviour among birds: this is, as one would expect, an adaptation to spatial and temporal variation in food sources. Such heterogeneity is also nearly ubiquitous: temporal variation in temperature or precipitation is present for a large fraction of avian habitats, both on a regular seasonal basis as well as on irregular shorter time-scales (Rappole 2013, p. 3). As a consequence of this basic fact, migration, in its various forms, is considered to be one of the oldest adaptations of the Aves class, and may be as old as flight itself (Rappole 2013, p. 3). Migration allows bird species to exploit variation in resources, and according to the pattern of variation there are different types of adaptive migratory behaviour (Rappole distinguishes between ten different types: Rappole 2013, p. 6). Not all of these adaptations

should be considered to be specific to the Aves class: as Rappole cautions (Rappole 2013, p. 7), the boundary between migration and movement is vague.

Yet, even though EH is ubiquitous, not all bird species are migratory. Avian adaptations to EH do not simply involve migratory behaviours gaining a selective advantage over resident behaviours in all environments. Because it is often possible for local populations to instantaneously exploit temporary increases in resources (the so-called “irregularity principle” of Willis: see Rappole 2013, p. 162), this leads to what is known as “resource superabundance” (Rappole 2013, p. 160), and allows for migratory species to enter habitats without competing directly with the resident species. So, the selective dynamics arising from EH can be viewed in two ways: as modifying the selective environment by creating a selection pressure for migratory behaviour (because migrants will be able to exploit the heterogeneity in resources), and as creating a new selective environment, by placing migrants in their own ecological niche so that they avoid selective competition with residents.

In general, this refers to a distinction to be made between two types of adaptive process: adaptation within a common selective environment, and adaptation by avoiding selective competition and exploiting novel ecological opportunities. In the first type, the EH-exploiting type gains a selective advantage over the EH-ignoring type; in the second type, the EH-exploiting type radiates away and does not selectively compete with the EH-ignoring type within a common selective environment. Besides migration, general motility, cooperation, and metabolic adaptations like endothermy are also assigned this adaptive role in the scientific literature.¹² A novel, un- or under-exploited pattern of EH can refer to the resources of unpopulated habitats, or to a spatiotemporal pattern in existing resources; in general it can refer to variation in any environmental variable that presents an ecological opportunity, including previously inaccessible resources, or patterns of predators, competitors, or potential co-operators.

While this process of adapting to novel ecological opportunities is not directly selective (rather, it is one of ecological radiation), it does involve the avoidance of selective competition. Patterns of EH outside the selective environment only become (relative) ecological opportunities when selective competition within the selective environment creates a dearth of opportunity. The latter follows from the basic fact that all resources are limited: this is not just the Malthusian rationale for natural selection within a common environment (Darwin 2008),

¹² See AUTHOR forthcoming.

but also the precondition for adaptive radiation as a way of escaping selective competition (Grant 2013; Tan et al. 2016).

The reason why this type of adaptive dynamic is relevant for the purposes of this paper is that novel, underexploited patterns of EH are ubiquitous. To begin with, they have been central to evolutionary history from its origins. For instance, if one considers one of the most viable hypothesised scenarios for early evolutionary history, namely that the last universal common ancestor inhabited hydrothermal environments (Weiss et al. 2016), this shows how concentrations of solar radiation were still untapped.

From a basic (thermodynamic) level of analysis, this is not a coincidence. A complete absence of heterogeneity in variables that affect fitness (including concentrations of energy, or concentrations of chemical elements) implies thermodynamic equilibrium, and this precludes metabolism and even the formation of cell boundaries. Physical heterogeneity is thus necessarily present in life-supporting environments. Moreover, of the 30 chemical elements used in biochemical reactions sustaining living organisms, only 4 are (currently) found in abundance: oxygen, nitrogen, carbon, and hydrogen. All the others occur patchily, and hence represented exploitable patterns of EH for early life forms.

Moreover, such novel patterns are also created by the evolution of life itself. For instance, the very presence of organisms represent an ecological opportunity for others, favouring the evolution of predation – which in fact happened early on, and predation is hypothesised as a selective driving force in the origin of the eukaryote cell and the origin of multicellularity (Pérez et al. 2016). The very metabolism of bacterial life creates steep gradients in the external environment through waste products – waste products which then may be exploited as metabolic inputs by other bacterial organisms (Fenchel 2002).

Hence, in sum, the way in which novel patterns of EH represent ecological opportunities is more than a contingent feature of natural environments: it was present at the origin of life, and as life evolved, more such ecological opportunities were created. Hence the process whereby EH-exploiting types alter their selective environments and colonise novel niches is not dependent on the contingencies of the geological succession of environments.

6. The Selectionist Rationale for a Trend in EH-exploitation

Recall the format of the selectionist rationale for a trend in M :

Given effectively random environments and without long-term constraints on natural selection, all evolutionary histories instantiate a trend in M that is caused by selection for increased M .

This section will detail the elements of the selectionist rationale for a trend in EH-exploitation: (1) the measure of EH-exploitation, (2) what the time-steps for EH-exploitation increases are, (3) what the expected trend would look like (under idealised conditions). The upshot is that at least one selectionist rationale for progress is grounded in the causal structure of natural selection. In this regard, I will make a tentative suggestion as to what concept of progress seems to be implied by the trend in EH-exploitation.

6.1 Measure of EH-exploitation. Since EH-exploitation refers to a type of interaction with the environment, the EH-exploitation of an organism (or lineage) can be spelled out in terms of the patterns to which the organism (or lineage) can react adaptively. Thus, the relevant lower-level space is not a space of possible genotypes or phenotypes, but rather a space of possible patterns that impact the fitness of organisms. This ‘EH-space’ has every indication of being *vast*: it contains all possible spatiotemporal patterns of variation in any primary resource or condition – including biotic variables. Finally, the ‘degree of EH-exploitation’ of an organism (or lineage) can be thought of as the volume of the region the organism (or lineage) occupies in EH-space¹³.

This measure induces at least a partial ordering: for at least some organism-pairs, a direct comparison can be made between degrees of EH-exploitation. Compare the sensory apparatus of cephalopod and bivalve species: bivalve ocelli track light intensity change (Gosling 2003, pp. 38–39) or at most can form low-contrast images (in scallop species: Speiser and Johnsen 2008), whereas cephalopod eyes can form high-contrast images (Boyle and Rodhouse 2005, p. 3); bivalves and cephalopods both have tactile sensors, but those of the latter are sensitive to surface texture (Boyle and Rodhouse 2005, pp. 25–26); through proprioception, hearing, or balance, cephalopods can pick up on environmental patterns to which bivalves are completely oblivious (Boyle and Rodhouse 2005, pp. 27–29).

However, the degree of EH-exploitation does not necessarily induce a total ordering. For instance, it does not seem that EH-exploitation in very different physical environments can be compared: a deep-sea angler fish is attuned to very different patterns than a European mole living underground. While both inhabit dark environments, there are obvious dissimilarities, and the question which of the two lineages possesses a greater degree of EH-exploitation is not

¹³ Note that, although independently derived, EH-exploitation seems to have numerous historical antecedents: “Perception of signals from the environment and ability to react to environmental stimuli”; “flexibility of behavior”; “ecological possibilities”; “control over the environment” (Rosslensbroich 2006)

likely to be a meaningful one. This by itself implies that any trend in EH-exploitation cannot be thought of as a linear increase in some scalar, as passive trends are often represented.

6.2 The Selectionist Rationale for EH-exploitation. Recall from the discussion of the neutral view that for a (passive) trend in measure M to occur, the probability of increase in M had an absolute size requirement (it cannot be too low) and a relative size requirement (it cannot be much lower than the probability of decrease in M). How do these requirements fare for EH-exploitation?

First, concerning the absolute size requirement: it is a certainty that living organisms inhabit environments with potential for further EH-exploitation. As previously discussed, opportunities for EH exploitation are present at the initial stages of an evolutionary history, and, moreover, these opportunities only increase as life evolves. This means that the absolute size requirement for a trend is met because novel and exploitable patterns of EH are present in every natural environment, albeit at varying spatiotemporal scales, and in different variables. A wealth of ecological strategies for increase in EH-exploitation present themselves: competition, predation, cooperation, and so on.

Second, the relative size requirement does not matter for a trend in EH-exploitation: there is no need to speculate about how probable the environment of a given lineage will favour EH-exploitation. The rationale for increased EH exploitation employs a different kind of thinking, one where the adaptive process alters the selective environment instead of selective environments exogenously driving the adaptive process. As detailed in section 6.2, increases in EH-exploitation involve exploring novel niches, thereby escaping from selective competition with EH-ignorers. Thus, a trend in EH-exploitation is a ratchet-like increase, involving radiation into new niches, rather than a trend caused by selective environments varying like coin flips.

This means that because evolving lineages change their own selective environments as well as those of other lineages, it no longer makes sense to evaluate the relative frequency of environments that favour increases in EH-exploitation versus those that favour decreases. This relative frequency is affected by the evolution of lineages: lineages constitute each other's environment. Specifically, even if whole clades were to face conditions that favour a decrease in EH-exploitation, this in itself creates novel opportunities for EH-exploitation in other lineages. Environmental heterogeneity is a multi-dimensional, consumable resource that cannot ever be entirely consumed: this is sufficient for a trend in EH-exploitation.

6.3 The Trend in EH-Exploitation. Selectionist rationales for trends are often thought to involve some scalar increase, whereby the later-evolved are more adapted or superior competitors compared to the earlier-evolved. By contrast, because the degree of EH-exploitation does not induce a total ordering, the implied trend in EH-exploitation is quite different. This selectionist rationale for EH-exploitation implies a radiation pattern of radiation that could be called *evolutionary unfolding*: over time, different parts of the totality of environmental heterogeneity are revealed to living organisms. At the origin of life only a few environmental variables (sources of energy) may be part of the selective environment, with the rest being part of the physical environment, but having no impact on fitness differences. As the totality of lineages ‘unfolds’ – i.e., increase the range of EH-exploitation – more environmental variables determine the selective environment of organisms. Biotic environmental variables come to represent novel dimensions of EH, and hence there is no upper limit to the degree of EH-exploitation: as life evolves, new opportunities for EH-exploitation are created.

Due to the ubiquity and necessity of novel exploitable patterns of EH, the trend is robust. Even an episode of a decrease in EH-exploitation across all lineages would subsequently create novel opportunities for other lineages. For instance, the decrease in motility in plants created new ecological opportunities for other lineages to graze on plants. The selectionist rationale for a trend in EH-exploitation remains unaffected by decreases in EH-exploitation.

Teasing out what the concrete empirical implications are of this view would be a separate task, beyond the scope of the discussion here. In any case, it can be said that tests for driven or passive trends (McShea 1994) would not necessarily be applicable. These tests, in attempting to measure the frequency with which lineages increase and decrease in M , assume that changes in M are induced by environmental changes that are independent of evolution itself. By contrast, in the pattern of unfolding, the evolution of a lineage affects both its and other lineages’ selective environment. Thus, an increase in EH-exploitation in a lineage may be followed by a long period of evolutionary stability; a decrease in one lineage may cause an increase in another lineage, and so on.

The main challenge for empirically testing a trend in EH-exploitation would be to develop auxiliary hypotheses, where an observable property is hypothesised to correlate with EH-exploitation in some way. Body size may be one candidate, since it is linked with increased motility and predatory capacity (Bonner 1988). Another candidate could be nestedness, insofar as the transitions to eukaryotes and multicellular organisms are thought to be associated with increases in motility, predatory capacity, and body size in EH-exploitation (Blackstone 2013;

Schirmer et al. 2011). Such auxiliary hypotheses would allow changes in EH-exploitation to be observed from the fossil record.

7. Discussion: Viable Progress Concepts

Progress concepts have typically been considered dead-on-arrival: regardless of whether or not ‘progress’ can be defined in a value-neutral way, they are not grounded in the causal structure of natural selection. In this paper, I have shown that the groundedness problem rests on a false premise about the role of environmental heterogeneity; when this premise is jettisoned, a very different picture emerges, whereby selection avoidance drives the exploitation of novel patterns of EH. Even though questions in evolutionary biology often turn on the issue of “how often does it occur?” rather than the issue of “can it occur?”, there seems no *prima facie* reason to believe increased EH-exploitation to be a marginal process, if nothing else because the process is associated with various fundamental ecological strategies – cooperation, motility, predation – that have had a major impact on evolutionary history.

The upshot is that the trend in EH-exploitation can be considered lawlike (in a way that seems to challenge Beatty’s ECT, although this will not be further explored here). Moreover, it shows the viability of at least two different progress concepts. As emphasised earlier, although such progress concepts may still be rejected for a variety of reasons – whether empirical or *a priori* – the point here is that, if they are to be rejected, this will not be because they are inconsistent with the causal structure of natural selection. Viability is understood relative to the nature of natural selection in heterogeneous environments.

The first progress concept is *unfolding progress*, referring to how some species or lineages succeed in exploiting patterns of EH. To revisit the example of cephalopods and molluscs: the former have better succeeded in exploiting EH than the latter. This does not mean they are better adapted in some narrow sense, since they inhabit distinct selective environments. Neither does it mean that the cephalopod has struck on one of a limited number of adaptive solutions to design problems. This is a concept of progress to be found in Conway Morris’ highlighting of convergent evolution (Conway Morris 2003, pp. 307ff) or Dennett’s notion of Good Moves in design space (Dennett 1995, p. 306). However, neither this concept of progress nor the naïve adaptive concept has a selectionist rationale that is grounded in the nature of natural selection. In unfolding progress, “better” is defined as a life form that can increasingly interact with and exploit a ubiquitous resource, even though the latter is at all times only partially accessible to extant organisms.

A second potentially viable progress concept is an eco-evolutionary version of *competitive progress*. Here “better” is defined as ‘more ecologically successful’, attributable to lineages that capture a larger share of energy in ecological networks than their ecological competitors. Ecological success and EH-exploitation are linked in the following way: while the capacity to exploit novel patterns of EH can lead to radiation, it can also allow for competitive (both intraspecific and interspecific) advantages, leading to a larger share of energy in foodwebs for some EH-exploiting lineages. For instance, motile organisms can seek ecological opportunity in a way that sessile organisms cannot. This is a concept of success often used (albeit rarely rigorously defined) by ecologically-minded scientists and philosophers, such as by Wilson (1990), Sterelny (2003), or Henrich (2016).

Since the measure of EH-exploitation does not induce a total ordering, it should not be a surprise that ‘ecological success’, as a proxy for EH-exploitation, does not allow for a total ordering either. There is little sense in comparing the ecological success of whole clades: the biomass of plants is, by latest estimations, 200 times that of animals (Bar-On et al. 2018); the Antarctic krill outweigh all wild mammals – let alone orcas – by at least an order of magnitude (Bar-On et al. 2018). Moreover, while EH-exploitation is linked to ecological success, it is clear that some clades have achieved certain ecological success without increased EH exploitation. For instance, sloths pursue a strategy of energy conservation; nonetheless, they account for a quarter of mammalian biomass in jungle habitats (Gilmore et al. 2001). Tardigrades have adapted to extreme environmental changes not by flexible avoidance but simply by resilience, by means of the mechanism of cryptobiosis (Møbjerg et al. 2011); yet, given their presence in every ecosystem (Møbjerg et al. 2011), tardigrades seem to have been ecologically successful to at least some degree.

In this way, a grounded selectionist rationale points to some potentially viable progress concepts. Even though the endorsement of progress-concepts – and of similar concepts, like ‘success’ – may be more widespread than commonly acknowledged (Ruse 1996), many will resist using this terminology. The aim of this paper was to merely to point to the flaws in the arguments that progress lacks theoretical underpinnings, and to give a positive account of how natural selection *could* provide such underpinnings.

8. Summary and Conclusion

Any selectionist rationale for progress must (1) be grounded in the nature of natural selection (and, in particular, should not rely on speculation about what types of selective environments are more probable than others), and (2) show how passive trends cannot always be more

parsimoniously explained as resulting from random walks. Even though the most common view on this matter is that both hurdles cannot be cleared, in this paper I have argued that the neutral view – which supports the sweeping generalisation that all passive trends can be explained as random walks – itself relies on speculation about the probabilities of selective environments. Then, based on the ubiquity of heterogeneity in natural environments, I proposed that selection for organismic capacities that exploit novel patterns of change in their local environment is a sound basis for a selectionist rationale. Since life itself creates patterns of change, there are always new patterns to be exploited.

The trend implied by this selectionist rationale involves changes in ecology first and foremost, rather than in morphology. Lineages radiate into novel ecological niches, and thus this trend where the totality of environmental heterogeneity (EH) unfolds looks nothing like the classic linear representations of progress. Life as a whole expands into niches where EH is exploited to greater degrees, without necessarily eliminating lifeforms that exploit EH to lesser degrees.

In showing that the hurdles facing selectionist rationales can be cleared, I do not pretend to have presented a full rehabilitation of the concept of evolutionary progress. However, what I do hope to have shown is that the search for a selectionist rationale for progress, grounded in natural selection, is eminently reasonable, and worthy of serious discussion – and not only in the context of science education, or the anthropology of folk biology.

[9732 words, incl. footnotes]

REFERENCES

- Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavioral and Brain Sciences*, 21(04).
<https://doi.org/10.1017/S0140525X98001277>
- Bar-On, Y. M., Phillips, R., & Milo, R. (2018). The biomass distribution on Earth. *Systems Biology*, 115(25), 6505–6511. <https://doi.org/10.1073/pnas.1711842115>

- Beatty, J. (1995). The Evolutionary Contingency Thesis. In G. Wolters & J. G. Lennox (Eds.), *Concepts, Theories, and Rationality in the Biological Sciences*. Konstanz: Universitätsverlag Konstanz - University of Pittsburgh Press.
- Berleman, J. E., & Kirby, J. R. (2009). Deciphering the hunting strategy of a bacterial wolfpack. *FEMS microbiology reviews*, 33(5), 942–957.
<https://doi.org/10.1111/j.1574-6976.2009.00185.x>
- Blackstone, N. W. (2013). Why did eukaryotes evolve only once? Genetic and energetic aspects of conflict and conflict mediation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1622), 20120266.
<https://doi.org/10.1098/rstb.2012.0266>
- Bonner, J. T. (1988). *The Evolution of Complexity by Means of Natural Selection*. Princeton University Press.
- Boyle, P., & Rodhouse, P. (2005). *Cephalopods: Ecology and Fisheries*. Oxford, UK: Wiley-Blackwell.
- Bradshaw, A. D. (1965). Evolutionary Significance of Phenotypic Plasticity in Plants. *Advances in Genetics*, 13, 115–155. [https://doi.org/10.1016/S0065-2660\(08\)60048-6](https://doi.org/10.1016/S0065-2660(08)60048-6)
- Calcott, B., & Sterelny, K. (2011). *The Major Transitions in Evolution Revisited*. MIT Press.
- Conway Morris, S. (2003). *Life's Solution: Inevitable Humans in a Lonely Universe*. Cambridge, UK: Cambridge University Press.
- Darwin, C. (2008). *On the origin of species* (Rev. ed.). New York: Oxford University Press.
- Dawkins, R. (1997). Human Chauvinism. *Evolution*, 51(3), 1015–1020.
<https://doi.org/10.1111/j.1558-5646.1997.tb03686.x>
- Dennett, D. C. (1995). *Darwin's Dangerous Idea*. London: Penguin Books.
- Dollo, L. (1893). The Laws of Evolution. *Bulletin de la Société belge de géologie, de paléontologie et d'hydrologie.*, 7, 164–166.

- Fenchel, T. (2002). Microbial Behavior in a Heterogeneous World. *Science*, 296(5570), 1068–1071. <https://doi.org/10.1126/science.1070118>
- Gilmore, D. P., Da Costa, C. P., & Duarte, D. P. F. (2001). Sloth biology: an update on their physiological ecology, behavior and role as vectors of arthropods and arboviruses. *Brazilian Journal of Medical and Biological Research*, 34(1), 9–25. <https://doi.org/10.1590/S0100-879X2001000100002>
- Godfrey-Smith, P. (2001). Three Kinds of Adaptationism. In S. H. Orzack & E. Sober (Eds.), *Adaptationism and Optimality* (pp. 335–357). Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511609084.012>
- Gosling, E. M. (2003). *Bivalve molluscs: biology, ecology, and culture*. Oxford ; Malden, MA: Fishing News Books.
- Gould, S. J. (1988). On replacing the idea of progress with an operational notion of directionality. In M. H. Nitecki (Ed.), *Evolutionary Progress* (pp. 319–338). Chicago: University of Chicago Press.
- Gould, S. J. (1989). *Wonderful Life: The Burgess Shale and the Nature of History*. W. W. Norton & Company.
- Gould, S. J. (1996). *Full House: The Spread of Excellence from Plato to Darwin*. Cambridge, MA and London, England: Harvard University Press.
- Gould, S. J. (2002). *The structure of evolutionary theory*. Cambridge, MA: Belknap Press of Harvard University Press.
- Grant, P. R. (2013). Adaptive Radiation. In J. B. Losos, D. A. Baum, D. J. Futuyma, H. E. Hoekstra, R. E. Lenski, A. J. Moore, et al. (Eds.), *The Princeton Guide to Evolution* (pp. 559–566). Princeton: Princeton University Press. <https://doi.org/10.1515/9781400848065-077>

- Grosberg, R. K., & Strathmann, R. R. (2007). The Evolution of Multicellularity: A Minor Major Transition? *Annual Review of Ecology, Evolution, and Systematics*, 38(1), 621–654. <https://doi.org/10.1146/annurev.ecolsys.36.102403.114735>
- Henrich, J. (2016). *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton: Princeton University Press.
- Hubbell, S. P. (2001). *The unified neutral theory of biodiversity and biogeography*. Princeton: Princeton University Press.
- Huneman, P. (2010). Topological explanations and robustness in biological sciences. *Synthese*, 177(2), 213–245. <https://doi.org/10.1007/s11229-010-9842-z>
- Johnson, N. A., Lahti, D. C., & Blumstein, D. T. (2012). Combating the Assumption of Evolutionary Progress: Lessons from the Decay and Loss of Traits. *Evolution: Education and Outreach*, 5(1), 128–138. <https://doi.org/10.1007/s12052-011-0381-y>
- Kauffman, S. A. (1993). *The Origins of Order: Self-organization and Selection in Evolution*. Oxford, UK: Oxford University Press.
- Kimura, M. (1968). Evolutionary Rate at the Molecular Level. *Nature*, 217(5129), 624–626. <https://doi.org/10.1038/217624a0>
- Lange, M. (2013). What Makes a Scientific Explanation Distinctively Mathematical? *The British Journal for the Philosophy of Science*, 64(3), 485–511. <https://doi.org/10.1093/bjps/axs012>
- Levins, R. (1968). *Evolution in Changing Environments: Some Theoretical Explorations*. Princeton, NJ: Princeton University Press.
- Louis, A. A. (2016). Contingency, convergence and hyper-astronomical numbers in biological evolution. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 58, 107–116. <https://doi.org/10.1016/j.shpsc.2015.12.014>

- Maclaurin, J., & Sterelny, K. (2008). *What Is Biodiversity?* University of Chicago Press.
- Maynard Smith, J., & Szathmary, E. (1995). *The Major Transitions in Evolution*. OUP Oxford.
- Mayr, E. (1994). The Resistance to Darwinism and the Misconceptions on which it was Based. In J. H. Campbell & J. W. Schopf (Eds.), *Creative Evolution?!* (pp. 35–47). Boston, MA: Jones & Bartlett Learning.
- McShea, D. W. (1994). Mechanisms of Large-Scale Evolutionary Trends. *Evolution*, 48(6), 1747–1763. <https://doi.org/10.1111/j.1558-5646.1994.tb02211.x>
- McShea, D. W. (1996). Perspective Metazoan Complexity and Evolution: Is There a Trend? *Evolution*, 50(2), 477–492. <https://doi.org/10.1111/j.1558-5646.1996.tb03861.x>
- McShea, D. W. (1998). Possible Largest-Scale Trends in Organismal Evolution: Eight “Live Hypotheses.” *Annual Review of Ecology and Systematics*, 29(1), 293–318. <https://doi.org/10.1146/annurev.ecolsys.29.1.293>
- McShea, D. W. (2001a). The hierarchical structure of organisms: a scale and documentation of a trend in the maximum. *Paleobiology*, 27(2), 405–423. [https://doi.org/10.1666/0094-8373\(2001\)027<0405:THSOOA>2.0.CO;2](https://doi.org/10.1666/0094-8373(2001)027<0405:THSOOA>2.0.CO;2)
- McShea, D. W. (2001b). The minor transitions in hierarchical evolution and the question of a directional bias. *Journal of Evolutionary Biology*, 14(3), 502–518. <https://doi.org/10.1046/j.1420-9101.2001.00283.x>
- McShea, D. W., & Brandon, R. N. (2010). *Biology’s first law: the tendency for diversity and complexity to increase in evolutionary systems*. Chicago ; London: University of Chicago Press.
- Møbjerg, N., Halberg, K. A., Jørgensen, A., Persson, D., Bjørn, M., Ramløv, H., & Kristensen, R. M. (2011). Survival in extreme environments on the current knowledge of adaptations in tardigrades. *Acta Physiologica*, 12.

- Payne, J. L., Boyer, A. G., Brown, J. H., Finnegan, S., Kowalewski, M., Krause, R. A., et al. (2009). Two-phase increase in the maximum size of life over 3.5 billion years reflects biological innovation and environmental opportunity. *Proceedings of the National Academy of Sciences*, *106*(1), 24–27. <https://doi.org/10.1073/pnas.0806314106>
- Pérez, J., Moraleda-Muñoz, A., Marcos-Torres, F. J., & Muñoz-Dorado, J. (2016). Bacterial predation: 75 years and counting! *Environmental Microbiology*, *18*(3), 766–779. <https://doi.org/10.1111/1462-2920.13171>
- Rappole, J. H. (2013). *The avian migrant: the biology of bird migration*. New York: Columbia University Press.
- Raup, D. M., & Gould, S. J. (1974). Stochastic Simulation and Evolution of Morphology—Towards a Nomothetic Paleontology. *Systematic Biology*, *23*(3), 305–322. <https://doi.org/10.1093/sysbio/23.3.305>
- Raup, D. M., Gould, S. J., Schopf, T. J. M., & Simberloff, D. S. (1973). Stochastic Models of Phylogeny and the Evolution of Diversity. *The Journal of Geology*, *81*(5), 525–542. <https://doi.org/10.1086/627905>
- Rosslenbroich, B. (2006). The Notion of Progress in Evolutionary Biology – The Unresolved Problem and an Empirical Suggestion. *Biology & Philosophy*, *21*(1), 41–70. <https://doi.org/10.1007/s10539-005-0957-4>
- Ruse, M. (1996). *Monad to man: the concept of progress in evolutionary biology*. Cambridge, Mass: Harvard University Press.
- Schirrmeyer, B. E., Antonelli, A., & Bagheri, H. C. (2011). The origin of multicellularity in cyanobacteria. *BMC Evolutionary Biology*, *11*(1), 45. <https://doi.org/10.1186/1471-2148-11-45>
- Sepkoski, D. (2016). “Replaying Life’s Tape”: Simulations, metaphors, and historicity in Stephen Jay Gould’s view of life. *Studies in History and Philosophy of Science Part*

- C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 58, 73–81. <https://doi.org/10.1016/j.shpsc.2015.12.009>
- Shanahan, T. (2001). *The Evolution of Darwinism: Selection, Adaptation, and Progress in Evolutionary Biology*. Cambridge University Press.
- Sober, E. (1984). *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. University of Chicago Press.
- Sober, E. (1994). Progress and Direction in Evolution. In J. O. Campbell & J. W. Schopf (Eds.), *Creative Evolution?!* (pp. 19–34). Boston, MA: Jones and Bartlett Publishers.
- Speiser, D. I., & Johnsen, S. (2008). Comparative Morphology of the Concave Mirror Eyes of Scallops (Pectinoidea)*. *American Malacological Bulletin*, 26(1/2), 27–33. <https://doi.org/10.4003/006.026.0204>
- Sterelny, K. (2003). *Thought in a Hostile World: The Evolution of Human Cognition*. Wiley.
- Tan, J., Slattery, M. R., Yang, X., & Jiang, L. (2016). Phylogenetic context determines the role of competition in adaptive radiation. *Proceedings of the Royal Society B: Biological Sciences*, 283(1833), 20160241. <https://doi.org/10.1098/rspb.2016.0241>
- Turner, D. D. (2011). *Paleontology: a philosophical introduction*. Cambridge ; New York: Cambridge University Press.
- Vermeij, G. J. (1987). *Evolution and Escalation: An Ecological History of Life*. Princeton University Press.
- Weiss, M. C., Sousa, F. L., Mrnjavac, N., Neukirchen, S., Roettger, M., Nelson-Sathi, S., & Martin, W. F. (2016). The physiology and habitat of the last universal common ancestor. *Nature Microbiology*, 1(9), 1–8. <https://doi.org/10.1038/nmicrobiol.2016.116>

- Werth, A. (2012). Avoiding the Pitfall of Progress and Associated Perils of Evolutionary Education. *Evolution: Education and Outreach*, 5(2), 249–265.
<https://doi.org/10.1007/s12052-012-0417-y>
- Williams, G. C. (2018). *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton University Press.
- Wilson, E. O. (1990). *Success and Dominance in Ecosystems: The Case of Social Insects*. Oldendorf/Luhe: The Ecology Institute. <https://www.int-res.com/articles/eebooks/eebook02.pdf>. Accessed 1 December 2018
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the sixth international congress of genetics* (pp. 356–366).