



This electronic thesis or dissertation has been downloaded from the University of Bristol Research Portal, <http://research-information.bristol.ac.uk>

Author:

Howard, Andrew S

Title:

Evolution, progress and biological function

a naturalist perspective

General rights

Access to the thesis is subject to the Creative Commons Attribution - NonCommercial-No Derivatives 4.0 International Public License. A copy of this may be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>. This license sets out your rights and the restrictions that apply to your access to the thesis so it is important you read this before proceeding.

Take down policy

Some pages of this thesis may have been removed for copyright restrictions prior to having it been deposited on the University of Bristol Research Portal. However, if you have discovered material within the thesis that you consider to be unlawful e.g. breaches of copyright (either yours or that of a third party) or any other law, including but not limited to those relating to patent, trademark, confidentiality, data protection, obscenity, defamation, libel, then please contact collections-metadata@bristol.ac.uk and include the following information in your message:

- Your contact details
- Bibliographic details for the item, including a URL
- An outline nature of the complaint

Your claim will be investigated and, where appropriate, the item in question will be removed from public view as soon as possible.

Evolution, Progress and Biological Function : A Naturalist Perspective.

Andrew Stuart Howard

A dissertation submitted to the University of Bristol in accordance with the requirements for the award of the degree of Doctor of Philosophy in the Faculty of Arts, School of Arts, Department of Philosophy.

Word count: 74,726

Originally submitted: 26th September 2019

Re-submitted with minor corrections: 13th January 2020

Page left intentionally blank

Abstract

The concepts of ‘progress’ and ‘evolution’ share a long and entangled history stretching back into antiquity. Pre-Darwin *all* evolution was thought to be progressive, now from our modern perspective *no* evolution is progressive; our commitment to naturalism appears to place the concept of progress out of biological reach. In this dissertation we show that this modern perspective is mistaken. Progress as a natural, biological concept is both attainable and is in fact widely used throughout evolutionary biology, albeit somewhat implicitly. Defenders of the ‘fact/value gap’ may be forgiven for some initial scepticism on this point, as could those concerned that the aim of applying ‘progress’ to biological phenomena is invariably motivated by anthropocentric wishful thinking; we make clear that evolutionary progress provides no grounds for human exceptionalism, and do so in a way which should not disturb David Hume or G.E.Moore.

Organisms *do* things: they eat, excrete, fly and photosynthesise - they have functions. These are quite unlike any properties belonging to non-living entities and phenomena outside of the biological world, and their singular nature warrants a singular explanatory approach. We rely heavily on the ‘Selected Effects’ conception of function to provide warrant for the special handling of biological function, but are far from doctrinaire; the means through which the SE conception warrants normative interpretation, is quite different than its supporters have supposed.

Evolution is thought to have produced many patterns in the natural world. Some are phantoms, the products of observers like ourselves hard-wired to seek regularity in disorder and intention and design in inanimate processes. Functions are not phantoms, they belong to the natural world, as does the progressive nature of the evolution produced by their incremental improvement.

Page left intentionally blank

Acknowledgements

Sincere and considerable thanks are due to my supervisors: Samir Okasha for his time, insight and patience, and his driving home the importance of recognising that Philosophy of Biology should respect both philosophy *and* biology; and Anthony Everett, whose guidance was never more than email away.

I would also like to thank Elselijn Kingma, whose seminars in the Philosophy of Medicine at KCL first drew my attention to the special problems that biological functions present for philosophy; and to Alex Hillman, for being a discerning and receptive sounding-board for ideas.

Lastly I would like thank Tim Ashton for his fine work in creating this dissertation's illustrations, the delays in this process were not his.

Page left intentionally blank

Dedication

This work is dedicated to my Mum, Beth for her support and encouragement, and to Pieter, without whom this dissertation might not have been completed.

Page left intentionally blank

Author's declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's *Regulations and Code of Practice for Research Degree Programmes* and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED:

DATE:

Table of contents

Chapter 1: Progressive evolution: introduction and overview

1.1:	Three objections	1
1.2:	Progress: intuitions and prohibitions	5
1.3:	Underlying themes	10
1.4:	Précis of the chapters ahead	11

Chapter 2: Natural selection: causality & creativity

2.1:	Introduction	18
2.2:	The non-casual argument	22
2.3:	Population-level causes and explanations	27
2.4:	The non-creative argument	39
2.5:	Lessons from the non-causal and non-creative positions	51
2.6:	Chapter recap and conclusions	58

Chapter 3: Normativity, function & progress

3.1:	Introduction	60
3.2:	Naturalising normativity & normativising nature	63
3.3:	Normativising nature #1: The ‘obligations’ of biological traits	72
3.4:	Normativising nature #2: The performances of biological traits	79
3.5:	Functional improvement and progressive evolution	87
3.6:	The prevalence of functional improvement	95
3.7:	Chapter recap and conclusions	98

Chapter 4: Natural patterns: authentic & quasi

4.1:	Introduction	103
4.2:	Patterns: elements, relations and mechanisms	106
4.3:	Authentic and quasi patterns: two real world examples	119
4.4:	Wilson and Dawkins: authentic or quasi?	123
4.5:	Higher-order mechanisms: some desiderata	135
4.6:	Chapter recap and conclusions	139

Chapter 5: Progressive evolution within lineages: representation & reality

5.1:	Introduction	142
5.2:	Delineation and anthropocentrism	145
5.3:	Demarcation: an A to Z of linear evolution	151
5.4:	The evolution of Cit+	164

5.5:	Nilsson and Pelger and the evolution of the eye	169
5.6:	Cross-species functional improvements	175
5.7:	Chapter recap and conclusions	180

Chapter 6: Evolution, progress and biological function: review and conclusions

6.1:	Three objections reappraised	182
6.2:	Four central themes revisited	187
6.3:	Direction, drive, improvement... and design	192
6.4:	Conclusion	197

List of tables and illustrations

fig 2.1	28
fig 2.2	43
table 2.1	44
table 2.2	48
fig 3.1	64
fig 3.2	92
fig 5.1	165
fig 5.2	170

Bibliography	200
--------------	-----

Page left intentionally blank

Chapter 1

Progressive evolution: introduction and overview

1.1 Three objections

Evaluative terms like ‘progress’ and ‘improvement’ were at one time as closely-associated with biological evolution as much as ‘natural selection’ and ‘adaptation’ are today. Naturalists preceding Darwin adhered closely to the origins of the term in considering ‘evolutionary’ changes as an ‘unfolding’ or ‘development’, when species *evolved* they *improved*. As with so much else in biology this changed after Darwin, although the diminishment of evolution’s progressive aspects was more gradual than might be supposed. The Malthusian picture of species competing in the struggle for survival with ‘natural selection’ in the role of arbiter seemed, if anything, to support rather than question evolution’s progressive credentials, driving organisms and species to ever greater strength and sophistication. Indeed Darwin himself at times struggled with it. Post-Darwin opinion initially remained divided on the issue, although from the advent of the twentieth century onwards the tide turned inexorably against an explicitly progressive interpretation. The growing confidence that Darwin was right, and that the theory of natural selection required no other supplementary forces to account for evolutionary phenomena was strengthened further by the identification of the mechanisms responsible for inheritance and the advances of the evolutionary synthesis made possible by this discovery. (Richards 1992, Ruse 1996, Shanahan 2004).

From the perspective of modern biology now able to manipulate DNA, map the human genome, and develop artificial genetic material, the concept of progress in relation to evolution would seem to belong to a different era, a perspective largely shared by contemporary philosophy. We may broadly categorise the reasons for this antipathy as belonging to three camps.

The first of these is scientific: knowing as we do the empirical facts of natural selection, and informed by our ever growing knowledge of genetics and ontogeny there is no place for the concept of progress - other than human *scientific* progress - concerning evolutionary biology. And further, no theoretical reason why we should think that the generation of random variation coupled with differential selection should produce, let alone amount to, advance or improvement. Evolution by natural selection can explain changes in the composition of populations, and changes in population-members' characteristics as a consequence, but this does not provide grounds to think that these changes embody progress in any way. Mature evolutionary theory it seems simply has no need for the concept.

The second of these camps could be construed as political: the abuse of evolutionary theory in attempting to justify the horrors of eugenics and racial 'purity' do not need to be spelt out here; we need only note that the distance between terms such 'progress' and 'improvement' and those more politically charged such as 'inferiority' and 'degeneracy' can at times appear distressingly small. Given this history - and indeed partial resurgence (Fuller 2018) - an aversion to using concepts which *rank* - rather than *sort* - the products of evolution is understandable. More recently, the welcome dismantling of the 'biological class system' has been more expansive. A rejection of biological superiority seems implicit in recent attempts to direct the concerns we have for ourselves towards the natural world and its non-human inhabitants more generally. The suggestion that some species could be superior *per se* to others- or treated as though they are- is increasingly recognised as untenable, environmentally damaging, and - in common with the notion of racial superiority - potentially giving a veneer of scientific respectability for abuses carried out by members of the putatively 'superior' category.

The third of these camps is philosophical: the recognition of biology as an empirical science naturalistically emancipated from supernatural characterisation was a battle hard fought and is a hill upon which many contemporary philosophers are happy to perish. Although it may look very like there is purpose, direction, intention and design in evolution, we must be explicit that the use of these terms is metaphorical - albeit perhaps unavoidable - when describing the causes of evolution and the phenomena that evolution can produce. Evolution

by natural selection is a natural mechanism, and the reintroduction of ‘improvement’ and ‘progress’ into the evolutionary lexicon is a regressive step in the wrong, anti-naturalist direction. The naturalist/anti-naturalist dialectic is one philosophical front, the naturalist/normative dialectic is another. The application of terms such as ‘progress’ and ‘improvement’ would seem to presuppose the existence of ‘standards’ or ‘values’ in order for any comparison of evolutionary phenomena to be carried out. But ‘standards’ as normally construed belong to people and to practices: rules of etiquette, safety limits in engineering and so on; appealing to the ‘standards’ of nature itself appears to be quite odd and unwarranted. How are these ‘standards’ set? How do we go about finding them out? And perhaps most importantly- ‘how is our understanding of nature improved by thinking of it in this way?’ If evolutionary biology is to be considered a respectable naturalistic, empirical enterprise its currency must be in facts; it should not be in evaluative concepts more at home in ethical or aesthetic discourse.

Yet despite these camps’ concerns the problem of progress in evolution refuses to go away. Not only is lay use of the term ‘evolution’ almost synonymous with ‘gradual improvement’, moreover the idea of evolution as a progressive, improving, process has been held firmly and explicitly by a number of leading evolutionary scientists: J.Huxley (1942), T.Dobzhansky (1956), B.Rensch (1959), F.Ayala (1974), G.G.Simpson (1974), E.Mayr (1988), D.Raup (in Nitecki (eds.) 1988), W.Wimsatt and J.C.Schank (in Nitecki (eds.) 1988), J.T.Bonner (1988), E.O.Wilson (1991), R.Dawkins (1996a), M.T.Ghiselin (1999), S.Conway-Morris (2003), and B.Rosslbroich (2014) for example. It is difficult to conclude that we should put these scientists’ collective opinion down to the lingering influence of 19th century metaphysics and move on.

Our aim in the work ahead will be to show that a wholesale rejection of progress in evolution is a mistake. Despite us enjoying respected - if select - biological company this positioning sets us firmly against the tide of philosophical orthodoxy. Crossing the well-recognised chasm between natural description and normative evaluation is normally fraught with philosophical peril. We hope to show that this tension can be diffused somewhat

concerning certain evolutionary phenomena, while simultaneously respecting the concerns from both the scientific and political camps.

If the scientific objection is correct and there is little reason to think that a mechanism based upon the biased retention and propagation of random variation should generate ‘improvement’ or ‘progress’ in anyway, there is equally no *prima facie* reason to think that the mechanism of natural selection should rule out the production of such phenomena. We do not usually demand that the empirical consequences of theories should be immediately derivable from their basic formulations. A line of argument based upon our inability to expect progress or improvement from natural selection could be just as easily applied to other common evolutionary phenomena such as speciation, sexual reproduction and so on; there is nothing in the basic tenets of natural selection that would lead us to expect these widespread phenomena either yet we hold natural selection to be the mechanism responsible. Even if natural selection does not predict ‘improvement’ or ‘progress’, as with these other evolutionary phenomena, this is a long way from saying that natural selection cannot account for them. Moreover the contention that we can expect change through the action of natural selection but not ‘improvement’ may not be as secure as is traditionally thought. As we will see, once we delve deeper into the reasons why certain variants provide the selection advantages they do the language of ‘progress’ and ‘improvement’ is both justified and empirically appropriate.

Concerning the political unease which treating evolution evaluatively in this way might arouse, our strategy in this dissertation will be to avoid this as much as possible by taking human evolution out of the equation as much as possible. We will establish our case for the restricted application of terms such as ‘better’, ‘improvement’, and ‘progress’ and so on in regards to non-human evolutionary phenomena first before specifically human evolution is introduced. Even then this will fall far short of a defence of human superiority. This should not be seen as ducking the issue; clearly *Homo Sapiens* are as much as product of natural selection as are any other species, but all too often philosophical approaches to evolutionary progress are compromised by anthropocentrism, either by deliberate inclusion of the author (Spencer 1891, Huxley 1942), or by providing a too-easy target for their critics (Simpson

1949, Gould 1988). By largely omitting human evolution from what follows our intention is to show that a progressive interpretation of evolution does not require any commitment to the former, while insulating our position from the dangers of the later.

The philosophical problems of justifying the application of non-naturalist or normative terms to natural phenomena are not as easily circumvented however - these will be dealt with directly. A recognition of progress in evolution neither requires a return to an intentional, teleological perspective rightly discarded by contemporary Darwinism, nor need it involve an unwarranted application of ethical or aesthetic terms better used to appraise the behaviour of moral agents. Although the progressive nature of evolution can be understood quite naturalistically without recourse to non-natural forces, the normative aspects of evolutionary 'progress' and 'improvement' are more subtle and accordingly require more careful philosophical treatment. As with naturalism itself, normativism comes in different varieties, and the normativism involved in recognising established standards within evolutionary phenomena is very different from that required in the recognition of ethical or aesthetic evaluation. There has been much confusion in this area, both from those who maintain an absolute moratorium on the use of evaluative terms in empirical science on the one hand, and from those who contend that by virtue of selection biological entities have functions that they are 'supposed' to perform on the other. Neither of these positions are correct. Analysis reveals a third way. There is no philosophical prohibition against the use of performative terms when describing certain biological phenomena, but this does not provide further warrant to think species, organisms or traits have 'obligations' in any non-metaphorical sense. Our reasons for thinking this will be made explicit.

1.2 Progress: intuitions and prohibitions

There are as many conceptions of evolutionary progress as there are authors on the subject, and this dissertation proves to be no exception. While it is difficult to find common ground concerning evolutionary progress in the work of authors as distinct as are G.G.Simpson and E.O.Wilson, or as philosophically diverse as are Herbert Spencer and David Raup, there are

general guidelines - central properties - that those sympathetic to progressive evolution implicitly respect.

The first of these is *directionality*. All evolution involves change, but only a subset of evolutionary change is directional in the sense that a series of evolutionary events could be reliably arranged into the correct temporal sequence absent knowledge of the actual order in which they occurred (Ayala, in Nitecki, (eds.) 1988). Directionality *per se* however does not amount to progress. As was repeatedly emphasised by Gould (1988b, 1990, 1996, 1997a) and affirmed with considerable rigour by Brandon and McShea (McShea 1998, McShea & Brandon 2010) the nature of evolutionary processes are such that directionality can be expected of them in ways very different from those which even a liberal interpretation could describe as ‘progressive’¹. Directionality is evidently necessary if progress is to be made but it is far from sufficient.

The second theme weaving through the progressive literature concerns the causes of evolution rather than its effects. Whatever various writers think evolutionary progress is, it cannot be accidental: there must be a force, influence or mechanism responsible for the directional evolution produced. So in addition to being directional, progress in evolution must also be *driven*. Early naturalists considered the causes of progressive evolution to be either supernatural in the form of divine intent, or due to as-yet-undiscovered forces or substances, although for modern conceptions of evolutionary progress natural selection is the near-exclusive causal mechanism². For Julian Huxley (1942) progress consists in the greater control of- and independence from- an organism’s environment; for J.Bonner (1988) progress can be witnessed through the positive correlation between size and anatomical

¹ The increases in both anatomical complexity and diversity between species more generally is probably the clearest example of a directional evolutionary trend produced incidentally. Reproduction is an inexact process and as a result differences should compound generation upon generation unless these effects are countered by opposing influences. Kimura’s ‘Neutral Theory of Evolution’ (Kimura 1968) use as a basis for calculating evolutionary time based on the regularity of random and selectively neutral variation is another.

² Why does natural selection not enjoy total exclusivity? Championed by Rosslenbroich (2009, 2014) the apparent evolutionary trend towards greater ‘autonomy’ in the sense of increased self-reliance and resilience to external, environmental influences is a plausible one - which Rosslenbroich suggests may be driven by factors other than natural selection. Rosslenbroich however leaves this as an open question, and as such we will not pursue this line of enquiry further here.

complexity, for Richard Dawkins progress consists in the “...accumulating number of features contributing towards whatever adaptation the lineage in question exemplifies.” (Dawkins 1996a, p1017). These are very different perspectives on what evolutionary progress consists in, but there is no disagreement here as to the identity of the driving force behind these diverse effects. For Huxley, greater control of- and independence from the environment is selectively beneficial, as are Bonner’s increases in size and anatomical complexity, whereas Dawkins’ increase in number of contributory features is explicitly adaptive.

The third theme is an effect: *improvement*. For an evolutionary series to be progressive the sequence must not only be directional and driven, it must also leave whatever is evolving in a ‘better’ or improved condition. Given the vagueness of this qualification, various authors have widely different perspectives on what improvement consists in. G.G.Simpson suggests increases in ‘awareness and reactivity’ to be one facet of evolutionary progress, and ‘internal fertilisation’ and ‘postnatal care’ to be others (Simpson 1949); for Wimsatt & Schank (in Nitecki, (eds.) 1988) increases in the amount of genetic information promote ‘...the possibilities for, and capability of producing and maintaining a larger variety of more complex adaptations through evolutionary time’ (ibid p233); for Bernhard Rensch, it is increases in anatomical complexity, rationality, and the refinement of the nervous system which provide progressively evolving organisms with their advantage (Rensch 1959).

From a naturalistic standpoint directionality is unproblematic, many natural processes are directional in the sense that earlier and later stages of the process could be placed in correct temporal order; direction *simpliciter* in evolution is commonplace. The existence of ‘drive’ in evolution is also comparatively uncontentious, as long as the nature of the force responsible remains naturalistic. As natural selection is held near-universally responsible for evolutionary progress by modern supporters this too remains unproblematic. It is with the third theme: improvement, that proponents of evolutionary progress appear to depart from scientific explanation. Most naturalists - and naturalistically inclined philosophers - would be happy to acknowledge the existence of directional trends in evolution and be more than

happy to consider natural selection as responsible, but recognising such trends as progress or improvement is often thought a step too far.

Our perspective on evolutionary progress follows closely in the spirit of those preceding by respecting each of these central properties. The type of evolution which we recognise as progressive is directional: in that given stages belonging to a sequence of progressive evolution we would with some confidence be able to place each stage in the correct order in which they occurred; and our type of progressive evolution is also driven: in that in common with nearly all other post-synthesis conceptions the mechanism responsible for its generation is natural selection. However the progress defended in this dissertation also explicitly charts improvements brought about by evolution - we do not shy away from the term. As we will see there is nothing in our use of 'improvement' to which evolutionary science might object; indeed in many cases it is difficult to see how empirical analysis might describe the phenomena *without* the use of performative terms.

Our guiding aim in the work ahead is to show that progress in evolution is real. It does not rely on subjective or questionable evaluations of natural phenomena, nor does it rely on empirical facts or philosophical concepts not already in general circulation. Moreover the type of progress which we defend in this dissertation conforms with both informed and intuitive use. Organisms possess traits which are real, and traits possess functions which are real, and the level at which traits perform their functions is not something imposed on them by external appraisal but is a quality of the traits themselves. Put simply some traits perform their functions *better* than do other traits. When a directional sequence of evolution features improvement in a given function and is driven by the selection advantages that each functionally-improved stage provides, the sequence will be progressive.

By recognising evolutionary progress in this way, we hope to avoid some common pitfalls to which previous attempts have been vulnerable.

Firstly, by recognising evolutionary progress as functional improvement, progress does not become too permissive. We do not claim that progress is a near-necessary consequence of

evolution as per Herbert Spencer's conception of progress, characterised as the inexorable creep of increasing complexity and differentiation (Spencer 1891), nor does the qualification of 'progress' become vacuous by overuse. Only a subset of evolutionary phenomena are produced by natural selection, and only a subset of the phenomena produced by natural selection are progressive.

However secondly, by recognising progress as functional improvement we do not characterise progress as too rarified or unusual. We might recognise this as a failing of Huxley's conception (Huxley 1941), if progress consists in 'greater independence from and control of the environment' it is difficult to see how this might apply to the evolution of organisms outside of the narrow category of metazoans. The possession of traits, and the possession of functions however can be appreciated quite broadly. Progress could be witnessed in the functional improvement of the traits of bacteria: such as the locomotive abilities granted by a mobile flagellum; the traits of archaea: such as ability of cannulae to exchange nutrients with proximate conspecifics; and the traits of metazoans: such as the visual acuity bestowed by the possession of functioning eyes.

Lastly, the conception of evolutionary progress as brought about through functional improvement is not impressionistic, nor is it based upon a vague or ill-defined notion of what progress might consist in. E.O. Wilson (1991) for example provides us with certain examples of 'progressive' events: the origin of eukaryotic organisms, the Cambrian explosion, to name but two, events which permitted increases in biodiversity, but Wilson provides no further rationale for thinking biodiversity itself should be progressive. It may *seem* that way but - as our work in chapter 4 will reveal - impressions are not enough. By contrast, should we know what the function of a trait is, we have good grounds for knowing what improvement in that trait's functioning would be, and a realistic expectation that we would be able to recognise and empirically measure 'better-functioning' if we saw it.

1.3 Underlying themes.

Weaving through this dissertation are a number of thematic strands. Some of these will be focused on quite directly, while others will be more peripheral. Perhaps the most evident is the explanatory scope of natural selection. Familiarity has perhaps made this fact less striking than it actually is: that such diversity and intricacy of ‘design’ as is found in the natural world formerly requiring the powers of a divine creator to be made intelligible might be accounted for using such a profoundly simple theory. We take a population of reproducing entities with different characteristics, ensure that these characteristics are heritable, ensure that differences in character causally influence reproduction, add an open-ended source of random variation, allow considerable time, and what began as a relatively simple population of abiotic entities incrementally evolves into the complexity and diversity of the living world that we witness today. Phenomena which formerly required intention and design are accounted for through the operation of a mindless and aimless naturalistic process. We will notice however that accounting for progress in evolution places additional and surprising demands on natural selection’s explanatory scope.

Our second subsidiary theme concerns the relationship between normative terms and natural phenomena. Given that natural selection is considered sufficient to account for such intricacy and ‘design’ it is perhaps surprising that the established position on what natural selection can and cannot explain should extend to: i) the existence of the traits and behaviours of organisms and species but stop short of: ii) warranting an evaluative understanding of the evolutionary processes which created them. Running through the work ahead is the thought that some aspects of nature are ‘normative’ independently of human interpretation, this warrants using normative terms to describe them.

Our third theme concerns the interplay between the representation and the reality of evolutionary phenomena. All representation to some extent *misrepresents* and this in turn may bias our understating of the phenomena represented. Representation may omit properties of the phenomena which it does possess or equally lend the phenomena qualities which it does not. Awareness of this can mitigate misunderstanding to some degree, but the

means of representing evolutionary phenomena: the metaphors, the intentional and teleological nature of explanation, the ‘historical’ methods used to interpret evolution, are pervasive and lie deep in evolutionary science. We remain vigilant to this throughout.

Our final underlying theme concerns the methodology used in understanding evolutionary phenomena. The usual approach for supporters of progress can be characterised as a ‘top-down’, pattern-centred approach. Previous attempts to characterise progress in evolution commence with the identification of ‘progressive’ phenomena: the biodiversity of the living world, the complexity of living organisms, the adaptive fit of species to their environments and so on - an impressionistic exercise. The work ahead makes full use of this approach, but not exclusively. A ‘bottom-up’, process-centred approach is also required. A pervasive thought in what follows is that the means of generation for such phenomena - its mechanism - should also be at the core of the concept of progress in evolution. This is usually thought of as expendable in progressive accounts but our thinking in the work ahead indicates this to have been a mistake.

We will later reappraise these themes in the light of the work to follow.

1.4 Précis of the chapters ahead

We start in **chapter 2** by defending the causal and creative powers allocated to natural selection in the face of recent philosophical attempts at reappraisal. Despite this being something of a rear-guard action in that we defend the powers of natural selection that orthodoxy already acknowledges it to possess, the analysis required proves to be instructive. There are two arguments here: one defending natural selection as a *casual* process, and another defending natural selection as a *creative* process, and both must prove successful if natural selection is to act as a basis for evolutionary progress that subsequent chapters require.

Questions concerning natural selection's causal potency are typically directed towards the extent of phenomena that natural selection can explain, or the nature of entities through which natural selection can be said to exert its causal influence. Opinion in both of these areas is divided. In terms of extent, discussion might be directed towards identifying the scale of natural selection's contribution for certain phenomena in contrast to the contribution of other evolutionary influences: drift, mutation/developmental factors, and migration. The long running debate concerning adaptationism is probably the clearest example of this type of inquiry: do we consider natural selection's role in evolutionary change to be such that these other factors can be safely marginalised in explaining the evolution of traits and characteristics? How big a part do chance events or contingent factors influence evolution in comparison to natural selection? Questions such as these presuppose natural selection's causal efficacy.

However, running alongside these debates concerning the extent of natural selection's causal powers, there is another which fundamentally questions whether natural selection has any causal efficacy at all. The position which we characterise as the 'non-causal view' considers the traditional picture of natural selection as a causal mechanism changing the composition of populations to be misconceived. The genuine site of causal interaction is between individuals and their environment, it is this and this alone which dictates the fate of individuals and their characteristics. Individuals interact, change in population composition is the effect; the traditional view of natural selection as a causal mechanism is misconceived, natural selection is the accumulated consequence of these individual interactions, not their cause. In response to this striking claim, we examine Roberta Millstein's suggestion that exclusive focus on the causal interactions of individuals is a mistake in evolutionary explanations. Populations of individuals have properties which individuals do not and some of these are causally efficacious in a way which cannot be explained using the non-causalist individual-centric perspective. We find Millstein's argument to be convincing, we notice however that the ground being fought on between the non-causalists and Millstein is somewhat different than their arguments might suggest. The causal efficacy of individuals versus the causal efficacy of populations arises due the conception of individuals and populations as distinct entities. This is incorrect. Individuals

are parts of populations, and populations are composed from individuals; the debate is not one of causality but of mereology.

As the name indicates, natural selection is a *selection* process, a point very much capitalised on by defenders of what we then characterise as the ‘non-creative view’. For nature to favour the propagation of certain types over others, these types must be different in some respect; natural selection requires differences first in order that they then be selected from. This again leads to quite a striking conclusion. If natural selection’s causal efficacy stretches only as far as favouring or disfavouring pre-existing variation, evolutionary science would appear to be mistaken by holding it responsible for the *creation* of complex and adaptive traits. The genuine site of biological creativity is to be found in the genomes of individual organisms and their development; mutation and genetic recombination create, natural selection promotes or eliminates. In contrast to the arguments of Millstein and the non-causalists, the ground on which the non-creative argument is fought *does* concern the nature of causality. We recognise that natural selection’s efficacy is limited to the favouring or disfavouring of pre-existing variation, but note that through this ability the background conditions necessary for the creation of complex, adaptive traits are produced.

With natural selection’s causality and creativity secured, our attention in **chapter 3** turns to the questions of how natural selection favours or disfavors certain types of variation, and how these empirical ‘facts on the ground’ should be understood philosophically. This chapter answers both the worries of the scientific and philosophical camps concerning progress in evolution. We begin by noticing that the reconciliation of natural facts and normative values typically attempts to explain the capacity for normative judgment, or the reasons for the judgments themselves, naturalistically. In contrast, when we describe evolution as ‘progressive’ and its effects as ‘improvements’ we are instead attempting to ‘normativise’ natural phenomena. We see that despite natural selection’s efficacy extending to the differential propagation of random variation only, the reasons for some variants having the selection advantages that they do is due to these variants performing their biological functions to an improved standard. This claim is supported using the ‘selected effects’ conception of biological function, and through doing so we recognise that

improvement in biological functioning is not a quirk or marginal characteristic of natural selection but instead is a fundamental property. We note that in opposition to the scientifically-oriented concerns that natural selection can explain change, but that it cannot explain improvement, a considerable amount of natural selection cannot be understood *without* recourse to such progressive terminology. Contra the philosophical worries that progress in evolution imports normative terms into scientific discourse without warrant, we suggest that what normativity is imported by use of evaluative concepts such as ‘improvement’ and ‘progress’ is fully justified by the nature of the phenomena described.

In **chapter 4** we widen our perspective from chapter 2’s defence of first principles, and chapter 3’s defence of normative interpretation, to appraise the patterns suggested to have been created by evolution at the largest of spatial and temporal scales. Here we recognise that even the most striking of patterns in the natural world can be produced incidentally. ‘Quasi patterns’ such as these can provide us with no further insight into nature’s operation. Those patterns for which a mechanism can be found however - ‘authentic patterns’ - can often prove informative and revelatory. Once we have secured the conceptual means through which quasi and authentic patterns may be discriminated, we return to the established literature regarding evolutionary progress, and apply our analysis to two high profile ‘progressive’ evolutionary patterns.

The first of these patterns, suggested by Richard Dawkins, is proved to be authentic. The gradual co-adaptation which Dawkins characterises as ‘arms-races’ between predator and prey coheres neatly with our preceding analysis- and is in fact a special case of the progress we describe in chapter 3. The pattern of escalation described by Dawkins is not accidental - in that there is a clear, reliable and well understood mechanism available to account for the observed pattern. We find that the co-adaptation characteristic of arms-races respects additional desiderata in that the mechanism responsible for the pattern is able to ‘screen-off’ lower level events underlying a pattern’s surface phenomena; is able to regulate which phenomena belong to the pattern and which do not; and finally the mechanism which generated the pattern not only accounts for the pattern’s ‘elements’ - those phenomena that together comprise the pattern - but also accounts for the arrangement of the elements in the

specific configuration provided. Wilson's pattern comprised from various 'landmark' events in the history of life on Earth fares less well in the light of this analysis. Lacking an obvious underlying mechanism to explain both its elements and their arrangement we note Wilson's pattern appears arbitrary and unreliable; certain stages are largely accidental, and consistency seems to necessitate the inclusion of other missing elements into the pattern. Additionally, certain stages force us to tackle the minutiae of their specific occurrences - the pattern presented by Wilson does not 'screen-off' lower level events in the manner required. We close chapter 4 by considering what properties large scale progressive patterns such as that suggested by Wilson would need to possess for us to consider them as authentic, and note that the only plausible casual mechanism currently available for their generation is natural selection.

In **chapter 5** the 'political' issue of anthropocentrism held at bay until now is confronted. The blows dealt by Stephen J. Gould to the possibility of a coherent and naturalist approach to progress in evolution continue to reverberate, but here we find that his dismissal of evolutionary progress as anthropocentric 'wishful thinking' to be unconvincing. We additionally find Gould's aversion to conceiving of evolution as a 'linear' process, transforming one species into another to be unwarranted. Evolutionary lineages are real and denying this can - and indeed has - led to persistent confusion. We do not however dismiss all of Gould's concerns. While lineages are real, they are also easy to misrepresent. We make explicit two procedures which are commonly used in the identification of lineages, and the methods used to represent their evolution, with emphasis on evolution which has occurred within certain lineages over large amounts of time. The first of these procedures is *delineation*: the process of picking out from the many splits and diversifications of the 'tree of life' the narrow sequence that connects a descendant population with its evolutionary predecessors. Approaches to evolutionary progress which perform this procedure usually consider the lineage identified as somehow more central or important in evolutionary history. We support Gould in recognising this as a mistake, but note that it is only the reification of a lineage which is unwarranted, not the coherence of the lineage itself.

The second of these procedures is *demarcation*: the nature of evolutionary lineages is such that they may be composed from a vast amount of individual organisms and can chart quite radical changes in morphology. A preparatory step in understanding these changes is to divide a chosen lineage into ‘stages’ in order to appropriately index the evolution which the lineage has undergone. Despite this procedure’s common implementation, we note that demarcating a lineage using impartial means may miss out much of what we expect an evolutionary explanation to provide on the one hand, while demarcating according to qualitative aspects of the evolutionary changes themselves can severely mislead on the other. Unlike the procedure of delineation which picks out a real and objective part of the natural world, we suggest that the procedure of demarcation will always be project specific; any one delineated evolutionary sequence may be demarcated in any number of different ways.

We look at three examples of demarcation: the first is one of our own, the evolutionary lineage linking the modern whale with its prehistoric land-based ancestor, *Pakicetus*; the two remaining examples are from evolutionary literature: the evolution of the novel phenotype ‘Cit+’ in the Long Term Evolutionary Experiment (LTEE); and evolution of the vertebrate eye in work of Nilsson & Pelger. We note that the way in which these lineages are demarcated greatly influences our attitude to the nature of their evolution.

Finally in **chapter 6**, our valedictory tasks are to recap, reappraise and conclude. With the benefit of our undertaken analysis we focus on the objections spelt out at the start of this chapter from empirical science, naturalist philosophy and those concerned that a defence of progress in evolution may open the door for ‘political’ abuse. We confirm that none of these camps should have grounds for suspicion. We then revisit this dissertation’s central themes: the explanatory scope of natural selection, the relationship between normative terms and natural phenomena, the relationship between representation and reality, and the methodology used in understanding evolutionary phenomena, appraise their influence, and realise that we have rarely stayed from at least one of them. Finally we turn our attentions to the properties traditionally thought to central to modern accounts of evolutionary progress: direction, drive and improvement. We recognise our perspective to be very much in keeping

with these former approaches, but additionally suggest that our analysis provides something which previous attempts do not- an empirical account for the appearance of design in the natural world as generated by progressive evolutionary processes.

Chapter 2

Natural selection: causality & creativity

2.1 Introduction

Perhaps the greatest success of natural selection is explaining the remarkable fit of organisms to the environments that they inhabit. The phenomena are familiar and ubiquitous: fish have gills and sleek fusiform bodies well-suited to their aquatic habitat, while cacti possess fleshy stems and sharp spines which aid water retention and deter thirsty herbivores in their arid surroundings; organisms are well-suited to their environments as a matter of course. A brief summary of how naturalists accounted for this phenomena might go as follows: stretching back into antiquity it was almost universally believed that organisms were expressly designed so as to suit their environments, until at the beginning of the 19th century it was gradually realised that organisms that changed could also account for their environmental suitability if they did so by adjusting their characteristics in response to the environment in appropriate ways. The mysterious mechanism responsible for achieving this was originally thought to be located within the organisms themselves, with progress accrued over an individual lifetime of environmental adjustment retained and passed on by inheritance; a suggestion which became increasingly untenable empirically. The means of environmental accommodation or 'adaptation' were then externalised to processes occurring outside of the individual organism; a development which signals to many the turning point at which evolution was naturalised (Ruse 1996, Bowler 2009). It was not necessary for individual organisms to change their own characteristics to suit the environment, given readily available conditions the environment would 'naturally-select' between existing character-variants and promote the propagation of those with a greater environmental suitability over those with less by virtue of their differential reproduction. Given other fairly uncontentious conditions 'Paley's question' of apparent design (Neander 1995, Razeto-Barry 2013) could be answered satisfactorily by known biological properties and dynamics; no designer or mysterious forces required.

In the terms of accounting for the fit of organisms to their environments this final picture has remained essentially unchanged although debate has ranged on a number of subsidiary issues: the relative importance of natural selection in comparison to other causal influences in producing the design-like characteristics of organisms (Gould & Lewontin 1979, Dennett 1995, Dawkins 1997, Godfrey-Smith 2001, Sterelny 2003); the conditions required for natural selection to create adaptive traits when it does in fact do so (Lewontin 1970, Endler 1986, Brandon 1999, Godfrey-Smith 2007); and whether natural selection exclusively produces adaptations for organisms alone or whether populations, species or higher taxa could themselves possess adaptations by virtue of similar dynamics operating at these 'higher' compositional levels (Grafen 1984, Wilson 1992, Sterelny 1996, Gardner & Grafen 2009). What natural selection can and cannot do is clearly fertile ground for debate. Recently however, shared assumptions underlying these debates have themselves become candidates for philosophical reappraisal.

Attacks have come from two fronts. Firstly, for natural selection to do anything at all it must be *causally efficacious* and this, it has been suggested, fundamentally misconstrues what natural selection actually is. We may be able to *explain* the fit of organisms to their environment through natural selection to a satisfactory degree, but employing the term in this way is in fact an instrumental convenience rather than an accurate representation of empirical reality. Natural selection is not the *cause* of adaptive phenomena- it is not a cause at all, according to a committed group of philosophers- but a consequence of lower-level events. It serves as a kind of explanatory shorthand through which we conceptually group the actual but disparate causes of adaptive phenomena in a regulated and tractable way (Rosenberg 1994, Walsh 2000, Walsh, Lewens & Ariew 2002, Lewens 2004, Matthen & Ariew 2009).

The second front attacks natural selection's *creativity*. The traditional view which regards natural selection as a process that creates the remarkable adaptive fit of organisms to their environments through the slow accumulation of favoured traits over others is, according to the non-creative view, mistaken. Biological novelty is ultimately sourced from mutation, not selection. Once variation has arisen *then* natural selection retains, ignores, or - as is almost

always the case- eliminates that variation. As nature can only select from options that are already available, regarding natural selection as the creative source of variation and by extension adaptation is mistakenly awarding credit in the wrong place. The traditional view of natural selection creating solutions to answer environmental demands is, it is claimed, based on a fundamental but pervasive mischaracterisation (Sober 1995, Stegmann 2010, see Razeto-Barry & Frick 2011 for comprehensive overview of the ‘non-creative’ approach).

Clearly if true either of these arguments would have profound consequences. If natural selection does not perform the creative work of the mysterious 19th Century forces wholesale in successively shaping organisms’ environmental suitability then what does? At best, the widely-held conception of natural selection as a producer and refiner of adaptive fit would have to be heavily revised at worst conceivably thrown out all together. Remarkably perhaps, recent work in this area appears to support an attitude that falls between revision and what even sober assessment would call revolution (Walsh 2000, Matten & Ariew 2009, Brunnander 2013). Rather than direct our efforts towards a detailed rebuttal of these various approaches, our aim in this chapter will be to focus on shared general shortcomings of both the ‘non-causal’ and ‘non-creative’ positions. As our task in the work ahead is to defend the conception of natural selection as cause of progressive evolution establishing its causal and creative efficacy is an important preliminary step.

Notice that despite their apparent similarity³ in perspective the argument for ‘non-causality’ and the argument for ‘non-creativity’ are incompatible, at least *prima facie*. Should we hold the argument for non-causality then natural selection has no causal efficacy, so the contention of the non-creative argument that natural selection’s influence is entirely destructive is moot- it has no causal influence in need of accurate characterisation. Conversely, if we are convinced by the argument for non-creativity and believe natural selection’s influence to be limited to eliminating disadvantageous variation, then it must have *some* casual efficacy even if its powers are only destructive. With our eye turned towards the wider task of interpreting the effects of natural selection, it appears we may

³ And indeed, in some cases shared supporters (Walsh 1998, 2000)

potentially spare ourselves some heavy lifting by countering the argument for non-creativity only. Showing, contra the non-creative argument, that natural selection creates, *ipso facto* shows that natural selection causes, so there should be no further requirement to directly engage with the arguments for non-causality. Although tempting, this short-cut will be avoided.

Firstly, the apparently *a priori* prohibition on what natural selection can or cannot ‘do’ if shown to be non-causal may well more flexible than it first appears. We might, for example, agree that although natural selection is not actually casual, it retains an explanatory power which stems from viewing its influence as creative not as destructive. Holding such a view of natural selection as a kind of ‘quasi-creative mechanism’ is a position which, as I see it, should not be automatically dismissed out of hand. Our second and more important motivation for avoiding the short-cut is that the reason for both the non-causal and non-creative arguments’ ultimate failure is additionally instructive. Addressing both arguments not only provides us with a broader conception of what sorts of phenomena natural selection can be called upon to justify (the *creative argument*), and illuminates the means through which natural selection goes about producing them (the *causal argument*), but also illustrates that both negative arguments against these positions are themselves based on shared assumptions. Non-causalists and non-creativists alike approach causation as an exclusive province of interactions at the ‘smallest’ and ‘closest’ possible level of analysis (Jackson & Pettit 1992). For non-causalists these interactions will be the life-history events of individual organisms, the idiosyncrasies of their births, endurance, propagation and deaths; for non-creativists it is the developmental effects caused by novel mutations in regards to their phenotypic consequences. Non-causalists believe that any causality discerned above the level of individual organism/environment interactions are a statistical abstraction, making natural selection explanatorily convenient but causally redundant; in a similar fashion supporters of the non-creative view believe that the creative work of producing phenotypic novelty is only ever carried-out ontogenetically, again no further contribution from natural selection is necessary. Through addressing both the non-causal and non-creative approaches we will see their shared methodological attempt to isolate *the* site of evolutionary causality is misconceived and erroneously mischaracterises how natural

selection causes and how it creates. The traditional view explains not because it is convenient or because it makes disparate events amenable statistically, it explains because it accurately describes the causal and creative powers of natural selection.

We will proceed as follows. In **2.2** we layout the non-causal position, characterising natural selection as an effect rather than a cause as more traditionally construed; in light of arguments put forward by Millstein, in **2.3** we find the proposals of the non-causal view to be unconvincing: exclusive focus on the undoubtedly causal interactions of individuals proves to be an unsuccessful strategy in light of the fact that individuals are parts of populations, not the causes of populations. In **2.4** we examine the non-creative position and also find its arguments to be unconvincing for similar reasons. Novel traits arise because of the interactions between genetic inheritance and ontogenetic development- the properties of individual organisms as the non-creative argument claims, but the reasons for an individual's genetic inheritance necessarily involve the selective dynamics of the populations of its ancestors; again we notice the wider perspective is wrongly overlooked in preference of events occurring 'at' the level of the individual. In **2.5** we will synthesise our observations from this chapter's previous sections and notice that although the cumulative selection, retention, and propagation of unlikely genetic combinations is a necessary step to explain complex adaptive features it is not sufficient to account for the 'fit' of organisms to the environments that they inhabit. For that we need to more closely examine the reasons for the fitness benefits of successive genetic combinations. Section **2.6** summarises the chapter's findings and establishes this dissertation's work ahead.

2.2 The non-causal argument

What we will collectively label the 'the non-causal view' can be seen developing through a series of papers which increasingly questioned the causal efficacy of natural selection in progressively bolder terms (Rosenberg 1994, Walsh 2000, Walsh, Lewens & Ariew 2002, Lewens 2004, Matthen & Ariew 2009). At their core, the idea that while it may be *explanatorily* productive to employ natural selection as though it were a causal influence,

the only genuinely *causal* factors responsible for the survival and reproductive performances of organisms are the proximate interactions that take place between the organisms themselves and their environments. When organisms and their traits interact with their environments systematically to promote their own proliferation we say that they are ‘naturally selected’ but this, these authors contend, is misconceived. ‘Natural selection’ is a high-level *product* created by the similarity of individual-to-environment interactions and not a casual *process* in itself.

This striking view is perhaps best illustrated in the following analogy from one of its leading proponents (Walsh 2000). We should consider the effects that natural selection appears to cause in populations featuring variation, heredity and differential reproduction as we would the shadow of an aeroplane crossing the ground as the aeroplane flies overhead. Like natural selection, the shadow has measurable properties (in the shadow’s case: size, speed, shape, location etc) and each successive state of the shadow appears, to the naive observer, to result from causal influences on the shadow’s former state, propelling the shadow forward in the direction of travel, perhaps also increasing/decreasing in size, brightness and so on. But we of course know better. Appearances are deceptive and the path of the shadow is not caused by the effects of influences working on former ‘shadow-states’, all of the shadow’s properties and its direction of travel are generated at a more fundamental level by the properties of the aeroplane and the its position relative to the sun, ground, cloud cover etc. By analogy the genuine causal explanation of how the genetic composition of a population moves from one state to another is not through any direct actions *of* natural selection, but through the combined and disparate causes behind each survival, death and reproduction of the population’s individual members, each migration and each mutation. When enough of the causes of survival, death and differential reproduction are of a similar kind and have a consistent effect over enough individual instances we combine them together into one- like the various properties of the aeroplane, sun etc combine to produce the shadow- and say that *the* cause of population change is natural selection. In actuality both the evolutionary change and the phenomena of natural selection are effects, jointly caused by the accumulation of events at the genuinely causal level of individual to environment interaction.

Sterelny and Kitcher in a passage near-ubiquitously cited in the non-causal literature suggest:

“In principle we could relate the biography of each organism in the population, explaining in full detail how it developed, reproduced and survived, just as we could track the motion of each molecule of a sample of gas. But evolutionary theory, like statistical mechanics, has no use for such a fine grain of description: the aim is to make clear the central tendencies in the history of evolving populations.” (Sterelny & Kitcher, 1988, p345)

Extending this perspective, leading non-causalists Walsh, Lewens & Ariew contend:

“If evolutionary theory is a theory of forces it isn’t a theory about the statistical structure of populations (and vice versa)..... The only genuine forces going on in evolution are those taking place at the level of individuals (or lower) and none of these (and no aggregate of these) can be identified with either selection or drift.” (Walsh, Lewens & Ariew, 2002, p453)

But reassure us that:

“... it should not be alarming that there is a theory that explains by appeal to population-level statistical properties as opposed to individual causal properties. The kinetic theory of gases is an obvious example of such a theory. The properties described by the phenomenological gas laws - temperature, pressure, volume - are, like trait frequencies, population-level properties. The nature and relation of these ensemble-level properties is explained in terms of a statistical property..... it is the *statistical* property..... that explains the changes in the ensemble-level properties if the gas” (Walsh, Lewens & Ariew, 2002, p463- italics in original)

The ontological line in the sand being drawn here by the non-causalists is between genuinely causal individual-level properties on the one hand, and non-causal, statistical, population-level properties on the other:

“...change in the structure of the population is explained and predicted by appeal to some statistical property, an *average* of individual propensities” (Walsh, Lewens & Ariew, 2002, p463- italics in original)

The salient term in this statement is ‘explained’: as in *merely* explained. The genuine *causes* of changes in population structure by contrast are *not* to be found in an average of individual trait-propensities, but in the responses of actual individual traits in their actual environments.

Employing statistically-abstracted properties such as averages: life expectancy, average wage etc. are of course indispensable when attempting to get a broader understanding of the collective behaviour of groups of individuals. We understand the dispersion of gas as caused by the properties of its individual component molecules but report on the gas’s behaviour in terms of the properties of the gas as a whole- it’s pressure and temperature; we acknowledge that rises in inflation are caused by the collective decisions of countless of individual economic agents, but explain them through abstracted collective-properties like ‘demand’ outstripping ‘supply’. This procedure itself is uncontroversial. The non-causalists’ point is well taken here: despite its undoubted utility, explaining the behaviour of collections of individuals using statistical properties is a quite different enterprise from identifying the underlying causes responsible for the behaviour of the collective’s individual members. This can be easily recognised by appreciating that the insights provided by statistical analysis are possible *because* of its ability to balance-out the idiosyncrasies of individual causes in order to survey the broader behaviour of collectives.

For non-causalists, individual organisms in populations- like molecules of gas in a container or consumers in an economy- are the fundamental units of casual interaction. Organisms, molecules and consumers interact and cause; populations, pressure and inflation are the consequences. Although convention regards ‘natural selection’ as a causal influence effecting the genetic composition of populations, this- like a lay observer interpreting entropy as a causal agency effecting the gas’s behaviour - is an explanatory convenience and

misconstrues causal reality. Populations, like economies, are abstractions. The causes of evolutionary change, like the causes of booms and recessions are not because ‘higher-level’ population-properties like trait-frequency or consumer confidence causally interact with other ‘higher-level’ population-level properties like ‘risk of predation’ or consumer debt; evolutionary change occurs when individual organisms within a population causally interact with their environments resulting in the individuals possessing certain traits reproducing more successfully than individuals which do not possess those traits.

When this occurs the representation of the trait in the population increases. The traditional *causal* interpretation of this is to say that in the environmental circumstances in which the interactions of the individuals had taken place, that the traits in question had been ‘naturally selected’; as though ‘natural selection’ were an additional causal influence acting at the level of the population in parallel with the lower-level proximate causes behind the survival and reproduction of the population’s individuals. According to the non-causal view, the ‘natural selection’ of traits in specific environments - the increase in their populational representation due to the systematic effects of their traits - is an *effect* caused solely by the interactions of the individual members possessing those traits. So when we say that trait *t* or trait *u* is naturally selected, any causality that this statement appears to award to natural selection ‘acting on’ *t* or *u* is merely metaphorical. ‘Natural selection’ is a statistical abstraction that we read into the causes of individuals’ survival and reproduction; albeit an instrumentality productive one. The only genuine causes of organisms’ survival and reproduction are the events of their life histories that actually caused those individuals to survive and reproduce.

Clearly, adopting this perspective could have alarming consequences; not least of which would be a fundamental reappraisal of large swathes of philosophical and perhaps even biological orthodoxy based upon natural selection’s status as a causal mechanism⁴. As such it might be tempting to dismiss the claims of the non-causalists out of hand by applying a kind of argument from incredulity. But suppose that we feel that these claims have some

⁴ The levels of selection problem is explicitly addressed in this revisionary way by Walsh (Walsh 2000). If natural selection is not a cause, then we do not need to be troubled about what hierarchical level: gene, organism, super-organism etc selection is actually acting on. According to Walsh natural selection is not acting at all.

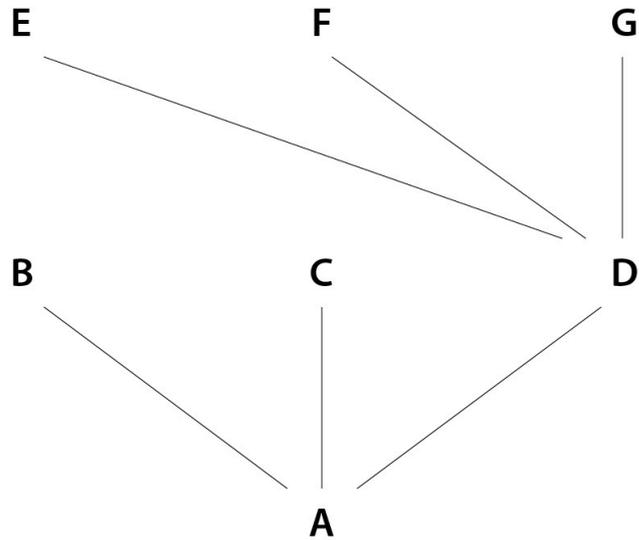
merit, but also suspect their call to revolution is somewhat premature. We seem to be presented with two options: i) contend that the causes of adaptive evolution are *not* exhausted by the interactions of individual organisms and conceive of natural selection as a casual influence distinct from these interactions, as the language of selective explanation appears to indicate- an approach which may strike us as equally avant garde as the claims of the non-causalists themselves; or ii) contend that the causes of adaptive evolution *are* exhausted by the interactions of individual organisms as per the non-causalists' claims, but that the causal influence of natural selection is included as part of these interactions; a project that can be characterised as one of *compatibility*. This will be our approach in 2.3.

2.3 Population-level causes and explanations

Our aim in this section is to demonstrate that in many situations, properties belonging exclusively to *populations* of individuals have an equal claim to casual efficacy as do the properties of the *individuals* involved. The main contention of the non-causalist position is that as causation happens exclusively to and through the interactions of individual organisms, any apparent causes not 'at' the level of individuals must therefore be merely statistical abstractions. This claim is central for the non-causalists, so their position would be fatally weakened if it were proved to be inaccurate. We do not deny individual-level causation by suggesting this - clearly individuals live, die and reproduce - a commitment to casual singularity does not also mean a commitment to the exclusivity of population-level causality, but neither does it mean a commitment to the exclusivity of individual-level causality as non-causalists contend. Evolutionary events can pertain of a plurality of *explanatory* perspectives, and we may reap the benefits of these without committing to a plurality of *causal* processes on the occasions that we do so.

Our case will be made primarily through an examination of proposals put forward by Millstein (2006) concerting population-level causal properties, but first we will consider the below simplified population (*fig 2.1*) to appreciate more clearly what the non-causalists are claiming.

fig 2.1:



We start with an individual A, which successfully produces three offspring: B, C & D; subsequently the D individual further reproduces three offspring: E, F & G, individuals B & C do not reproduce. For the non-causalists the causes for D's successful survival and reproduction are to be found exclusively in D's birth, encounters with predators, consumption of food and water and so on: events quite idiosyncratic to D alone. As we have a full account of the reasons for D's survival and reproduction provided with this information, the non-causalists contend, we have no need to introduce 'natural selection' as a causal mechanism additional to the facts provided; we already know the full casual story. So what is natural selection? The non-causalists' answer is this: if there were other members of the wider population (not shown in *fig 2.1*) who shared D's characteristics and who successfully passed on these characteristics in similar environmental circumstances to those in which those characteristics benefited D then we would say that those characteristics were naturally selected. In circumstances like this where similarity in inheritable characteristics has similar effects on survival and reproduction, it looks as though 'natural selection' is 'favouring' certain traits and 'disfavouring' others, but natural selection is not *causal* here. The individual's environmental interactions are causal; 'being selected' is a statistical

artefact, an *effect* brought about through repeated similarity in characteristic-to-environment reproductive success. We could, if we wished, exhaustively explain both the life histories of individuals, and the the development of populations made up from such individuals by not referring to natural selection at all.

Our first point of departure from this individual-centric perspective is to observe that given realistic ecological circumstances, even the very simplified and limited population structure featured in *fig 2.1* could have casual relevance for the depicted outcome. The compromised survival and reproduction of B and C may not have occurred in isolation to the life of D and may have had a profound causal influence on D's survival and reproduction; so when considering D's reproductive successes they cannot automatically be marginalised. The non-causalists are correct that D's life history provides us with the reasons for D's survival and reproduction, but this is somewhat myopic. A comprehensive explanation for the circumstances of D's life history would have to encompass ecological information on D's habitat, and that may depend significantly on the number and activities of the other members of D's *population*. For example: suppose the habitat of the above population is desperately short of resources making the carrying capacity of the population severely limited. It might well be the case that in order for D to successfully reproduce B and C did *not* reproduce. If B or C had survived to reproduce the resources for D would have been even more severely restricted, perhaps to the extent that it would have prevented D surviving long enough to successfully reproduce as it did. These are not contrived or artificial circumstances, the life histories of individuals within a population intersect as matter of routine.

We may be confident that the non-causalists' response to this criticism would be to agree that B, C and D's individual life histories may have significant overlap, but the correct way to interpret this is as individual life histories with individual-level causes intersecting and resulting in the involved individuals' survival and reproduction. But can notice here that this version of events is not quite right. D's reproductive success is not conditional specifically on the absence of the *individuals* B and C, it is conditional on there being no or very few other *population-members* in the second generation using up environmental resources at D's

expense. It just so happens that in the above population the required resources were available because B and C specifically did not survive to reproduce, but had there been other members of the population D's reproductive success could have been very different. D's comparative success was due to the population being small - a population-level property.

As population-size is clearly a property of a population, we might start to suspect here that at least some of the disagreement between the individual-level-causes of the non-causalists, and the population-level-causes of Millstein might hinge on differences in explanatory approach rather than differences in casual reality. For the non-causalists B and C's influence on D's reproductive output is conceived as the *life histories of two individuals intersecting the life history of a third individual*; in contrast for those holding a population-level casual view it is *the fact that population-size is very low* that explains D's reproductive success. It seems clear that these differences concern the mode of explanation, and do not concern differences in casual circumstances: B and C's deaths *means* that the population was much reduced, the population was much reduced *because* both B and C died. So even at this preliminary stage, we might suspect that evolutionary causality *per se* is perhaps not as central to the debate as the non-causalists contend. We will notice this theme repeated as we turn to the arguments of Millstein.

Populations of individuals of course have many properties that their members taken individually⁵ do not: distribution, diversity and ecological niche being obvious examples. Millstein's approach is very much in contrast to that of the non-causalists when she contends that some of these population-level properties are causally responsible for evolutionary effects that cannot be accounted for from an individual-level perspective. The first of these population-level causal properties is variation. Variation in population

⁵ As Millstein correctly points out, we must be careful not to conflate the proposals of the non-causal view with the 'levels of selection' problem as more commonly construed (Dawkins 1975, Sober & Wilson 1994, Sterelny & Kitcher 1998, Okasha 2006). It is not the primacy of the gene or of the organism etc for natural selection that is our concern here as that question presupposes the causal efficacy of natural selection; the very thing we are setting out to investigate. 'Individuals' for our purposes could mean individual genes or individual organisms, and 'populations' could mean populations of genes, or populations of organisms. Although the default non-causal position does in fact see the individual organism as the exclusive site of evolutionary causality (Walsh 2000, Lewens 2004), this is not- I think- an absolute requirement of the non-causal position.

members is a property of populations- not individuals- and has significant consequences for the differential reproduction of population members: itself a population-level effect.

That populations are comprised from individuals with various traits, that these traits have a direct bearing on the individuals' survival and reproduction is undeniable. The non-causalists are again correct on this point: if we were to take Sterelny & Kitcher's advice and trace the life-history of each individual in a population we could know in principle exactly how the traits of each individual contributed to their survival and reproduction. But there is much that looking at an individual's life history and reproductive output in isolation like this will *not* tell us. We would not know if that individual's interactions with its environment were systematic of individuals with their characteristics, and whether the number of offspring produced by that individual was typical of similar individuals or entirely idiosyncratic; to know that we need to know the responses of multiple individuals. Additionally, we would also not know how to interpret the number of offspring produced by an individual from a 'selection perspective' as this can only be understood by knowing *population-level* information concerning the reproductive outputs of other individuals.

At first pass this suggestion might strike us as quite odd. For surely it is individuals themselves that reproduce at various rates and 'number of offspring' is a property of individuals, not a property of populations; increase in population size being a product of its individual members' reproductive output. But consider again *fig 2.1*: this time we will stipulate that it depicts a very small section of a much wider population composed from substantially more individuals. Notice that we can do this without making *any* changes to the individual life-histories of the population-members shown. The fates of A, B, C, D and so on all may all play out exactly as originally specified, only this time there are other population-members present but remote enough not to have made any survival or reproductive differences to the individuals explicitly represented. Individual A successfully reproduces three times as does D just as before, whereas B and C do not reproduce at all, just as before. But the number of offspring alone (an individual-level property) cannot tell us if these individuals are selectively *favoured*. To know that we need to know the comparative reproductive outputs of the other population members. If the *entire* population

was as depicted in *fig 2.1* as originally specified, by virtue of having three offspring, A and D would be selectively favoured, whereas both B and C would not. But this is only because A and D's contributions to their subsequent generations are disproportional to the *population* average. If the average reproductive output of individuals in an extended population were to be above three offspring then this would mean A and D were *not* selectively-favoured as their proportional contribution to the next generation would be much-reduced. In clear breach of the non-causal view's marginalisation of the population-level, we can only appreciate the evolutionary significance of A and D's survival and reproduction by comparative information not necessarily provided by an analysis of the minutiae of individual life histories, no matter how forensically this is carried out.

We might again preempt the non-causalist response to this. Being 'selectively favoured' or 'not-favoured' are the statistically-abstracted consequences of the actual survival and reproduction of individuals with certain characteristics. They are not the causes of differential reproduction, but an optional characterisation of the results of genuinely casual, individual-level interactions. The very fact that the population dynamic in *fig 2.1* can be variously interpreted as demonstrating both positive and negative selection depending on the dynamics of the wider population- despite other population members having no casual relevance in either scenario- demonstrates that natural selection itself cannot be casual, but is instead a matter of context and interpretation. As such, it would seem that this example supports the non-causal position rather than argues against it. Millstein's contention that intra-population variation *causes* differences in reproductive success might seem to be in trouble at this point. Non-causalists can remain entrenched in their original position: it is the characteristics of D and the environment inhabited by D that caused D to survive and reproduce as it did, and it is the characteristics of B and of C and their interactions that similarity decided their reproductive fates. Although differences in characteristics directly effect differences in reproductive output, it is not the variation itself that is directly causal, but the specific characteristics and interactions of the individuals themselves.

Millstein's second, and more convincing, population-level causal property is genotype-frequency. Populations have genotype-frequencies whereas individuals have genotypes. The

population-level property of genotype frequencies causes the population-level effect of changes in genotype frequency. Although individuals may develop quite dramatically over their life-cycles (we may think of metamorphoses as a striking example of this) their genotypes do not. When a population features n copies of a certain genotype at time t_1 , and m copies of that genotype at t_2 , (and $n \neq m$) this is not due to any changes in the composition of the individuals themselves. Individuals die, reproduce and migrate, and the frequencies of genotypes in a population change when individuals with different genotypes undergo these processes at differential rates. We may characterise Millstein's second claim then as follows: under certain circumstances, the frequency of genotypes in a population (a population-level property) causes changes in genotype frequency over subsequent generations (a population-level evolutionary effect).

Consider the familiar hawk/dove population dynamic. A population 'T' features two genotypes 'd' & 'h'. Genotype d produces 'dove' behaviour: faced with aggression from other population members d-types always back down to avoid potential injury. In contrast h-type 'hawk' individuals will always follow-through on any threat despite coming out worse from such encounters some of the time. We can explain the wax and wane of population T's d and h genotype frequencies over subsequent generations as follows: as the frequency of d-types in the population rises, the lower the probability that any d-type will confront an aggressive h-type, so the lower the probability that behaviour characteristic of the d-type will be necessary to avoid potential harm. At some point d-type behaviour is no longer beneficial compared to h-type behaviour; there is a much-reduced possibility that any h-type will be met with a h-type response so their characteristic aggressive behaviour begins to pay-off as h-types also profit from the lack of h-types in the population, causing an increase in h-type representation at the expense of the representation of d-types. At some further point however, h-type behaviour ceases to be beneficial; the increasing risk of damage to h-types from other h-types as they become more frequent means that d-types regain the advantage and grow in populational representation.

If this should be considered an inaccurate *explanation* for the ongoing development of this population's genotypic frequencies, then the non-causal view is more revolutionary than our

earlier concerns might suggest. Models of frequency-dependence are entirely routine in evolutionary biology (Sterns 2000, Nowak & Sigmund 2004, Ridley 2004, Otto et al. 2008), so Millstein cannot be accused of capitalising on a rare phenomenon to make an unwarranted general claim about the casual efficacy of population-level properties. But as the preceding case study indicted, denying natural selection's explanatory relevance is not the non-causal view's contention:

“...natural selection theory *explains* changes in the structure of a population, but not by appeal to the individual-level *causes* of births, deaths, and reproductions. As natural selection theory does not deal in the causes of these individual-level phenomena, it cannot yield *causal* explanations for them” (Walsh, Lewens & Ariew, 2000, p469- italics added).

Their point is that explanatory adequacy and genuine causality can come apart, and that explanatory accounts which award natural selection casual efficacy achieve this *in lieu* of finer-grained explanations exhaustively detailing the life histories of individual population members. It is generally accepted that we may enjoy explanatory plurality while adhering to causal singularity (Jackson & Pettit 1992, Sterelny 1996); different explanations may provide different perspectives of certain events and answer to different types of enquiries, so this claim of the non-causalists is not in itself that contentious. But their suggestion implies something further. That out of the different available explanations some in particular will be more legitimate than others because they detail events occurring *at* the finely-grained level where causality is taking place: the interactions of individuals. In contrast however, as Millstein argues, it seems natural here for us to consider the relative frequencies of d-types and h-types - population-level properties - as *causally* responsible for the changes in genotype-representation over subsequent generations and not just as an explanatory substitute *in lieu* of finely-detailed information. It seems hard to avoid the conviction that any explanation for the fluctuation of genotype frequencies over generational time - and we may allow there could be a plurality of explanations - would have to include the frequencies of the genotypes themselves, and that the reason for their inclusion is not one of descriptive preference or perspective, but because they are directly causal to the evolutionary outcome that we are trying to explain.

Notice again here that in common with our population in *fig 2.1* there is no disagreement about the nature of cause: we know this is due to d-types' and h-types' interactions; the crux of the matter appears to stem from the non-causalists' commitment that causation occurs exclusively to a single individual. This is a different and much stronger claim, and one that we - and we may suspect, holders of the traditional view more generally - have good reasons to reject. When individuals within a population causally interact, then *ipso facto* the population also casually interacts, and when a population interacts it does so *through* the activities of its individuals. The cause is not in contention here. As we saw above, we may *explain* these interactions through the causal interactions of individuals (the deaths of B and C) or the causal interactions the population (very low population size); but considering 'the individual' and 'the population' as two *casually*-distinct entities and stipulating that we prioritise one as 'the' level of causality and dismissing phenomena not occurring 'at' that level as derivative is misconceived. This can be appreciated quite clearly by recognising that the actual relationship that holds between the interactions of individual organisms and populations of individuals is in fact quite different to the casual one that Walsh's analogy of a plane casting its shadow over the landscape suggests. The correct relationship of individual-to-population is not one of causality but of mereology. Aeroplanes cast shadows that are ontologically distinct from their physical causes and are at a spatial and temporal remove from them; individuals do not cause, and are not temporally or spatially separate from, populations and population-level events, individuals are *parts* of populations, affecting one *means* affecting the other.

The ontological 'line in the sand' drawn by the non-causalists between genuine causal properties on the one hand and statistical 'abstract' properties is an important one for the non-causalists, but it is somewhat artificial. Certain properties are able to be *both* statistical in the sense that they are derivable from the central tendencies of 'lower level' individuals *and* genuinely casual in that they describe the sufficient conditions required for the collective behaviour of the 'lower level' individuals to produce a given effect.

Consider again the analogy of pressurised gas heated within a container. As the gas heats, it expands and exerts increasing pressure against the walls of the container as it does so. At some point, the pressure of the gas surpasses the structural integrity of the container causing the container to crack, releasing gas and decreasing the pressure within. How should we understand this? Clearly we *explain* the effect here - the expansion of the gas - in terms of the properties of the gas: pressure, heat, and volume are not properties possessed any individual molecule. The irreducibility of these properties, coupled with their explanatory and predictive abilities would be, we might suggest, sufficient for them to be considered as directly causal. In contrast, the non-causalists would seem committed to the view that the ‘population-level’ properties of the gas, its heat and pressure, are not to be considered directly causal but statistical and derivative- as would the effect of the gas’s expansion. For the non-causal view the container’s cracking would be caused when one particular molecule struck with sufficient force.

Although the approaches here diverge on explanatory approach, it is clear that they pertain to the same underlying causal process and as such it is also increasingly clear that disagreement here is not over the nature of the causes themselves. The non-causal strategy is to dial the ‘grain’ of their explanation to that of the individual taking the actions of surrounding individuals as background conditions for the ‘main events’ of individuals’ survival and reproduction. In contrast, Millstein’s strategy is to set the ‘grain’ of her explanation higher to the level of the population, because only from this perspective can we discern that properties of populations like the frequencies of d- and h-types produce effects such as changes in their frequencies over subsequent generations; she considers the grain set by the non-causalists as too narrow to appreciate this. As we can see from the gas example, the non-causal view in claiming that their grain is the only one ‘at’ which casualty takes place is somewhat myopic. The effect of the behaviour of a single individual with other individuals seen as background conditions simply *is* the joint effect of all of those individuals when taken as a collective. There is no contradiction here for the non-causalists to capitalise on; it is just the explanatory emphasis which each side places on the relevance of the background conditions which seems to be in dispute.

Non-causalists maintain that the events surrounding the actions of an individual (here the relative frequencies of each genotype) should be considered as background, whereas Millstein considers the activities of all the collected individuals to be of equivalent causal importance. With this in mind, we might consider Millstein's claim against the non-causalists that causality is to be found exclusively 'at' the level of populations to somewhat less than a decisive win, but notice that although these individual-level and population-level explanatory approaches are compatible - indeed once it has been established that they describe different aspects of the same casual process, it is difficult to understand how they might not be compatible - the appropriate approach will be guided by the nature of the phenomena that we are hoping to better understand. We understand the reasons non-causalists believe a population-level phenomenon like natural selection to be derivative is that it cannot be discerned at the level of the individual; individuals cause individual-level effects which are then aggregated- in the non-causalist view- misleadingly. Similarly, Millstein considers population-level properties like the frequencies of d- and h- types to be directly, not derivatively casual; but the subsequent effects that she explains by using these population-level properties are population-level effects. If we wish to explain *individual*-level effects; births, deaths etc then, as-per the non-causal view, it seems we must construe causes as acting on or through *individuals*, but the converse of this also seems to hold: if we wish to explain *population*-level effects then it seems we must construe their causes as acting on or through *populations*. The suitability of explanatory approach, at least in these cases, appears to be orientated according to the level of the explanandum.

The principle in play here would seem to be quite a general one and not particular to specifically biological populations and individuals. We cannot explain why any particular molecule of gas in the pressurised container behaved in the way that it did using 'higher grain' properties such as pressure and temperature; for that we would require considerably more 'lower grain' information about that particle's specific location, speed etc and the specific conditions of the other particles⁶. Similarly with the actions of an individual

⁶ Notice the non-causalists' point is not entirely otiose: if we had access to this information we might omit the 'population-level' properties of pressure and temperature as a casual influence for the movements of any particular particle. This however is quite different from claiming that pressure and temperature are not causal. Cracks in containers can be *caused* by increases in pressure and burns can be *caused* by the temperature of pan handles, as much as differences in a population's genetic composition can be *caused* by natural selection.

economic agent: if we want to know why they acted in the specific way that they did: their choice of purchases, investments and so on, citing ‘inflation’ as a cause for their distinctive behaviour will be insufficient, we need to know the the ‘facts on the ground’ and options open to that particular agent, their financial acumen, ethical concerns and so on. This observation may not be of primary importance for Millstein whose objective is to show the causal efficacy of natural selection *per se* and which her use of frequency-dependence sufficiently demonstrates, but we require natural selection not just to be casually effectatious but to be responsible for specific phenomena: the uncanny fit of organisms for the environments which they inhabit, and there are reasons to doubt that changes in genotype frequency over subsequent generations alone are sufficient to achieve this.

The difficulty arises when we ask at which ‘grain’ our explanandum: the fit of organisms to the specifics of their environments, is to be interpreted. Should we stipulate that it is the fit of *individual* organisms to their environments, then as this appears to be an individual-level effect it would seem to require an individual-level cause. As Millstein’s example demonstrates only that natural selection causes population-level effects, we must either reformulate our original question so that the properties of individuals are construed as population-level properties, or ‘shift grains’ and explain the properties of individual organisms employing individual-level principles. Either of these approaches might strike us as unsatisfying. In the light of the our findings in this section, our key question would seem to hinge on whether we consider: *the fact that populations contain individuals which fit their environments more suitably than other individuals*, and: *the fact that individuals have certain features which fit their environments* to be equivalent explananda. And this would seem to depend on the nature of the features that we are interested in. Natural selection may explain why the representation of one genotype’s populational representation varies in comparison to the representation of another but it does not explain how these processes conspire to create the adaptive phenomena which prompted our original investigation. Recognising natural selection as a mechanism for generating evolutionary effects like changes in genotype frequency falls very short of our intended expectation for it as the mechanism responsible for the creation and promotion of complex adaptive traits. We will

return to this problem in greater length in 2.5, but first we need to establish natural selection's *creativity*.

2.4 The non-creative argument

Soon after natural selection's debut it was realised that without a source of phenotypic novelty the process of selection would eventually run out of things to select from (Huxley 1942, Vorzimmer 1963, Dobzhansky 1970, Mayr 1982); a puzzle that found its resolution decades later with the discovery that heredity was transferred by means of discrete units and that the copying processes of these units could, on occasion, be inexact. Genes could change, or mutate (Bowler 1978, Futuyma 2015), these 'mistakes' and their phenotypic consequences were able to account for the 'open-ended' source of novelty that the theory required.

As with the non-causal position, there is no single proponent of the 'non-creative' position (Endler 1986, Muller & Wagner 1991, Sober 1993 & 1995, Walsh 1998, Brunnander 2013) although again the core of the argument can be distilled quite straightforwardly. As per the original worries, natural selection cannot claim to be a *source* of evolutionary novelty or variation in its own right, but works exclusively by selecting *pre-existing* variation. The ultimate source of this variation in evolution is mutation, with phenotypic novelty producible through the recombination of existing genes. Without mutation, evolution by natural selection would eventually eliminate all fitness differences and grind to a halt, leaving a static and selectively-homologous population. But when naturalists assert that natural selection explains how and why complex adaptations like gills, leaves, lungs and tracheae *originated* their statements are not intended to be taken elliptically but as a description of empirical reality:

“Natural selection... is much more than a negative process... Natural selection may be compared rather to a painter which creates a picture by mixing and distributing pigments in

various ways over the canvas. The canvas and the pigments are not created by the artist but the painting is.” (Ayala 1970, p5)

“For most professional biologists and many thinking people in the 20th Century (both theists and atheists), the theory of evolution provides an adequate basis for understanding both life’s diversity and the adaptedness of living things. We believe evolution through natural selection to be the main creative force that adapts organisms to their environments, without comprehension or design.” (Doolittle 1994, p47)

Are these and other similarly-minded authors wrong to think this- should they be exclusively thanking the mechanisms of mutation and inheritance instead? In common with the *prima facie* claims of the non-causal approach we might be tempted here to think the position too outlandish to warrant serious attention. The claims of the non-creative position however are more subtle. Consider: for natural selection to occur in a population it must necessarily feature heredity, differential reproduction and *variation*, so as variation is an integral part of natural selection, and mutation is the ultimate source of variation, variation is therefore an integral part of natural selection also. The creative- and non-creative positions are in complete agreement in this. Additionally no one (or at least, no one who recognises natural selection as a causal mechanism) disputes that natural selection provides a causal explanation for the distribution of a trait through a population once that trait already exists. The difference in the positions becomes evident when we turn to the issue of whether the process of natural selection has creative potency *in addition* to the genetic variation that it requires. For holders of the non-creative position, mutation is the ultimate source of evolutionary novelty- with the proviso that genetic recombination may act as a secondary developmental mechanism:

“The processes of mutation, recombination, heredity and ontogeny determine which traits an individual has. A distinct kind of causal explanation- developmental explanation- appeals to these processes to explain of individuals, severally, why they have the traits they have.” (Walsh 1998, p252)

Natural selection on the other hand refers exclusively to:

“... the differential survival and reproduction of individuals. Natural selection sorts the variants, preserving those traits and combinations of traits which confer on an individual the capacity to survive and prosper in its environment. But it does not produce the variants, combine them, or ‘transmit’ them from parents to offspring. Selection merely causes changes in trait frequencies within a population. If that’s what it does, then that’s what it explains”. (Walsh 1998, p252).

Notice that in common with the non-causal argument, the non-creative argument overlooks events occurring ‘above’ the level of individual organisms- but in a different way. For non-causalists the fates of other population members were conceived of as background conditions for the site of genuine *causality*- located ‘at’ the level of the individual: its birth, death, reproduction etc. Here the mechanisms for genuine *creation* are likewise identified as those which occur to individuals: the specifics of how their genetic inheritance coupled with developmental factors are able to produce phenotypic novelty. In contrast to the non-causal position, natural selection is conceived as having *some* influence through the ‘sorting’ of traits, this is however quite distinct from its traditional characterisation as a mechanism responsible for the ‘production’ of variants promoting environmental/phenotype suitability.

Notice also that the structure of the non-creative argument is considerably more straightforward and less contrived than that of the non-causal argument. According to the non-creative view natural selection can never be creative *in principle*. Natural selection ‘sorts’ and selects, it does not ‘produce’ or create and so cannot be held responsible for creating adaptive traits. The non-creative argument stands or falls according to the accuracy of this condition, and as such provides a defence of the traditional view with a clear objective. If it can be demonstrated that the mechanism of natural selection can in certain circumstances show *additionally creativity* with no further mutational influence the non-creative argument fails.

The non-creative argument may be summarised as follows:

- i) Natural selection is (at minimum) a two-stage process. Stage 1 (S1) : the creation of variation, followed by Stage 2 (S2) : the interaction of that variation with the environment.
- ii) S1, the 'creative stage' of the process is uniquely sourced by genetic variation: these are mutation with recombination as a secondary mechanism.
- iii) S2, the 'interactive stage' of the process is natural selection's exclusive province. Here, some of the variants created at S1 interact with their environment so as to promote or impede their own reproduction.
- iv) Some variants have no effects in a given environment, but variants that do interact with their environments as described in iii) either increase their representation or reduce their representation.

Therefore:

- v) Natural selection as characterised as interaction in iii) either works to maintain, promote, reduce or eliminate the representation of a novel variant. Natural selection does not create new variants, so natural selection does not create adaptive traits which are comprised from many such variants.

Our approach will be to demonstrate that this argument is unsound due to premise ii) being untrue. Genetic variation is *not* the only source of evolutionary novelty, under easily-achieved circumstances S2 (the interactive stage) may itself produce evolutionary novelty and on the occasions that it does so the mechanism of natural selection is responsible. This defence will provide more than simple support for one side of an established dialectic; additionally, it will show that the traits that we are particularly interested in- complex adaptations, such as wings and leaves and so on which grant organisms their environmental suitability- are the ones that require the mechanism of natural selection to explain their existence the *most*. Mutation, recombination and other developmental factors are able to

generate phenotypic novelty and variation, but are insufficient in themselves to generate the specific *type* of accumulated and complex variation required to produce the traits that we find particularly compelling. In order for mutation and other developmental factors to account for these they require suitable conditioning and natural selection is currently the only mechanism that is capable of providing this.

We may consider the below phylogenetic diagram (*fig 2.2*: based on those found in Sober 1995 and Walsh 1998) as representative of the non-creative view's attitude towards the roles of genetic variation and natural selection.

Fig 2.2

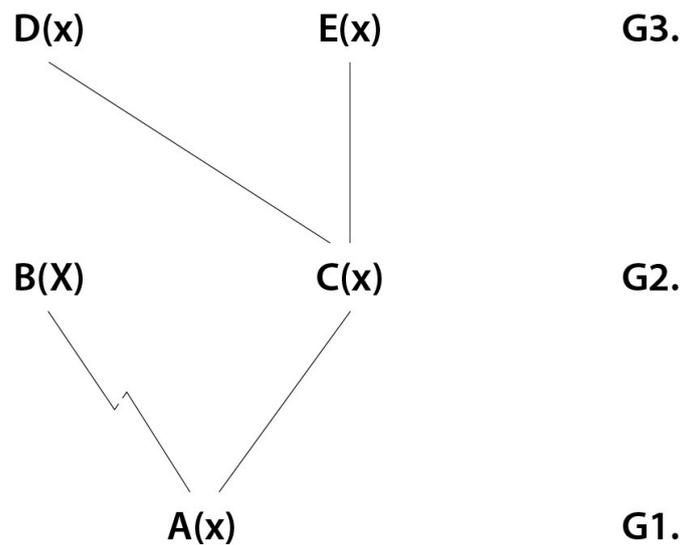


Fig 2.2 shows the evolution of a population (P1) over three generations: G1, G2 & G3; comprising five organisms: A, B, C, D & E; and two alleles: x & X. G1 is composed from A only; G2 composed from its offspring B and C; and G3 is composed from C's offspring D and E. The broken line between A(x) and B(X) shows a mutation event; here the mutation of the x allele into the X allele. x is selectively superior to trait X. The construal of natural selection as an exclusively destructive agent in accordance with the non-creative view can be shown in the following table:

table 2.1

P1	Variable: Mutation	Variable: Selection	Result: No. of Variants in population.
G1	-	-	1 (x)
G1 - G2	yes (x to X)	no	2 (x, X)
G2 - G3	no	yes (x > X)	1 (x)

We start with a single variant only in G1 (x). Then due to the effects of mutation we have two variants in G2, doubling the amount of variation. But in G3, due to the effects of selection the number of variants is reduced back to one only. In terms of creativity and destructiveness it seems fairly clear what roles genetic variation and selection are playing here. Had there been no genetic variation between G1 and G2 the number of variants would have remained at one, if there been no selection against variant X in G2 the number of variants would have remained at two. Mutation adds, selection removes. But this simplified perspective begins to breakdown when the population dynamics become more complicated; and, we may add, more realistic.

This time our population initially features two genes: x and y, each of which has one allele: x mutates into X, y mutates into Y. This time traits are produced by genetic combinations, rather than a single gene as in P1. We therefore have four possible traits: xy, Xy, xY, XY. Also distinct from the traits featured in P1 their fitnesses in this example will be malleable as our aim will be to asses the alternate projections for a population when its evolution is shaped by the relative fitnesses of the traits which comprise it, keeping all other variables fixed. These alternative dynamics will play out in two populations: Population 2- with selection, 'P2S'; and Population 2- with selection much reduced: 'P2N'. Our argument here is quite straightforward: if the *only* difference between P2S and P2N is the strength of selection, then we are justified in holding natural selection responsible for any variances in evolutionary outcome that come about. Moreover, if one of the differences in these populations is the creation of a novel trait that the other lacks, then we are also justified-

contra the non-creative position- in considering natural selection to be a creative agent under these circumstances and those appropriately similar.

We will establish the fitness scheme featuring selection in P2S as ‘FSS’:

FSS: $XY > xY = Xy > xy$.

In FSS, the XY composite trait provides a greater selection advantage than both the xY and Xy composites whose fitnesses are equivalent, but superior to composite xy.

P2S is initially composed from xy organisms exclusively. So for the novel composite trait XY to be created, the following steps need to be accomplished first: {x mutates into X, y mutates into Y}, then {both X and Y must exist in the same organism}. This second step could occur through genetic recombination in a sexual population or by a lineage inheriting either X or Y and proceeding to mutate the second required variant. As the population is initially composed from xy, any appearance of composite traits xY (when y mutates into Y) or Xy (when x mutates into X) can be straightforwardly explained by referring to actions of mutation alone. But because Xy & xY are selectively superior to xy according to FSS, we may expect natural selection to increase their representation in P2S once they have arrived provided that genetic recombination does not prevent inheritance by splitting their component letters up. Notice that introducing traits based on genetic combinations has already made the dynamic in P2S quite different from the dynamic in P1: there natural selection was responsible for *removing* the newly-mutated variant; here selection potentially maintains and *propagates* the genetic variant, and developmental factors- should the components of Xy or xY have been joined by genetic recombination- have a high probably of *removing* the novel variation.

In contrast, for our alternative population: P2N, we keep all variables and initial conditions the same as those in P2S aside from the relative fitness of each composite trait. We will establish the fitness scheme with much-reduced selection in P2N as ‘FSN’:

FSN: $XY > xY = Xy = xy$.

In FSN, as in FSS, the XY composite trait provides a greater selection advantage than the xY, Xy and xy composites, however in P2N the fitnesses for these composites are equivalent.

For the XY composite to arrive in P2N the same two mutational steps: {x mutates into X, y mutates into Y} then {both X and Y must exist in the same organism}, must occur just as they did in P2S, but here in P2N because xY and Yx provide no fitness benefit relative to xy in FSN, their initial spread through the population is far from assured. Relative to population size it becomes increasingly likely that both Y and X would be lost rather than drift to fixation absent the fitness dynamic provided in FSS. This clearly has severe consequences for the likelihood of P2N developing the novel XY trait⁷.

Our case for the creativity of natural selection can be set out as follows. When the trait in question is a composite trait (here XY), and there is positive selection for each step of the process between initial state and the final composite (here {x mutates into X, y mutates into Y} then {both X and Y exist in the same organism}) natural selection greatly *increases the probability* of the composite-trait's creation due to its ability to 'hold' each step in the population until the next step comes along. So, we while agree with the non-creative view that:

$$P(x | S) = P(x | \neg S)$$

The probability of a single copy of any given trait based on a single gene (here x) existing given selection (S) is equal to the probability of that trait existing without selection. Natural selection does not influence the probability of any one trait-creating mutation occurring. However, in opposition to the non-creative view:

⁷ The probability of a novel selectively-neutral mutation reaching fixation equals $1/(2N)$. In a population of 10 this works out to .05; in a population of 100 this works out to .005.

$$P(XY | S) \gg P(XY | \neg S)$$

The probability of a given composite-trait (here XY) existing *with* selection is much higher than the probability of that composite-trait existing without a selective process- given a fitness scheme like that described in FSS. The creativity of the selection process is that it makes biological features arising from *combinations* of traits vastly more probable, while having no influence on the mutations that create the original traits themselves.

The thinking behind this argument is itself far from novel. R.A. Fisher suggested that:

“... it was Darwin's chief contribution, not only to Biology but to the whole of natural science, to have brought to light a process by which contingencies a priori improbable, are given, in the process of time, an increasing probability, until it is their non-occurrence rather than their occurrence which becomes highly improbable” (R.A.Fisher 1954, p314).

And similar approaches characterising natural selection as a probabilistic cause of unlikely genetic combinations are hard to avoid: Kimura (1961), Neander (1988), Forber (2005) & Razeto-Barry & Frick (2011); as well as being implicit in the portrait of natural selection's creativity that Ayala (1970) paints in the metaphor as quoted above.

A table again helps makes the contrast between otherwise identical populations featuring different selection dynamics more explicit:

Table 2.2

	Variable: Mutation	Variable: Selection	Result: Composite traits in population	Result: Total number of composite traits created
P2N FSN: $XY > xY = Xy = xy$				
Stage 1			x1 (xy)	x1 (xy)
Stage 1 to 2	yes (x to X)	no	x2 (xy, Xy)	x2 (xy, Xy)
Stage 2 to 3	yes (y to Y)	no (Xy lost to drift)	x2 (xy, xY)	x3 (xy, Xy, xY)
Stage 3 to 4	no	no (xY lost to drift)	x1 (xy)	x3 (xy, Xy, xY)
P2S FSS: $XY > xY = Xy > xy$				
Stage 1			x1 (xy)	x1 (xy)
Stage 1 to 2	yes (x to X)	no	x2 (xy, Xy)	x2 (xy, Xy)
Stage 2 to 3	yes (y to Y)	yes ($Xy > xy$)	x2 (Xy, XY)	x3(xy, Xy, xY)
Stage 3 to 4	no	yes ($XY > Xy$)	x1 (XY)	x4 (xy, Xy, xY, XY)

The layout of *table 2.2* differs from *table 2.1* in that we measure changes through ‘stages’ rather than generations (which would be unworkably large); the first ‘result’ column shows the composite-traits present in the population at any given stage, while the second shows the total number of composite traits created.

Notice that the mutation order is fixed in both P2N & P2S: x to X between stage 1 to 2, y to Y between stage 2 to 3, no mutation stages 3 to 4. The single difference is the fitness differences as detailed in FSN & FSS. We clearly see in P2S that the effect of selection produces results very different from those shown in both P1 and P2N and that this influence is clearly creative. In P1 selection acted to remove the novel trait and return the population to its original state, in P2N it is the *absence* of selection that returns the population to its original state, while in P2S selection’s effect is to retain the novel intermediate enabling the later *creation* of XY. In P2N the total number of composite traits created without selection is x3, in P2S the total number of traits created with selection is x4: all those from P2N plus the additional XY composite, thus demonstrating natural selection’s creative potency.

This is the creative argument. Does it do what is required of it? Our aim was to demonstrate that premise ii) of the non-creative argument was false; to do that we had to show that genetic variation was not the only source of evolutionary novelty, and this, plausibly, could still be a matter of debate. Isn't natural selection just shuffling around what is already present? Mutation, not selection, is the reason x changes into X ; mutation, not selection is the reason y changes into Y ; and genetic recombination, not selection is the reason both of these develop jointly in an individual organism. Isn't natural selection's role in the creation of XY restricted to making it more *likely* that X and Y will conjoin rather than providing the causal means for their combination? In essence, the strength of our argument reduces to how liberal a conception of causality is considered appropriate for evolutionary phenomena. This was not an issue for the causal-argument, where the metaphysical nature of causation itself was left unexamined without consequence⁸. Typically though, as it is widely accepted that the type of causality which natural selection trades in is of the probabilistic variety (Kimura 1955, Beatty 1984, Yoshimura & Shields 1987, Mayr 1988, Glymour 2001, Orr 2005, Skipper & Millstein 2005) we may rightly feel the explanatory burden lies with the non-creativists in demonstrating why a different causal principle should apply here when evolutionary science itself is satisfied with a probabilistic explanation.

If we feel that the presence of natural selection, like mutation, should be an absolute *necessity* for the creation of a novel trait such as XY then an account based on *probabilistic* causality will naturally fail. XY *could* be created in the examples above without the assistance of natural selection; although it is vanishingly unlikely that this would happen due to high probability that by lacking superior fitness either of its components would be lost to drift soon after their arrival. Alternatively, if we feel that making events more probable is the appropriate causal standpoint for evolutionary phenomena then the creative position succeeds: natural selection is not necessary for the creation of XY (as this could happen incidentally), and natural selection is not sufficient for the arrival of XY (as superior fitness in both Xy and xY traits does not guarantee their proliferation), but natural selection does undoubtably make the creation of XY considerably more likely and this, we may

⁸ Although Millstein is at pains to point out that her argument for the casual potency of population-level properties is compatible with various different casual accounts (Millstein 2006, p8).

suggest, is sufficient to explain that trait's *creation*; and thus demonstrates natural selection's creativity.

We will close this section with a an observation. Notice that the lower the value of $P(XY | \neg S)$ the stronger the case that occurrences of XY require natural selection to create the appropriate circumstances for its creation. If the probability of any one selectively neutral variant reaching fixation is $P = 1/(2N)$, then the probability of any two neutral variants reaching fixation is $P = (1/(2N))^2$, three variants: $P = (1/(2N))^3$, and so on. Clearly the *higher* the number of variants required to make a trait, the lower the probability that mutation alone could produce it; and the larger the effective population size, the lower the probability that once a neutral variant had been created that it would proliferate rather than disappear.

At the start of this chapter our aim was to provide a defence of the traditional- casual and creative- view of natural selection as a mechanism responsible for the remarkable suitability of organisms for their environments. We have seen how natural selection has no effect (other than possibly eliminative) for the creation of single-gene traits which are due to genetic variation alone, but can have a significant effect on bringing separate variants together in a population when each is selectively beneficial, and the *greater* the number of variants that are required to produce any given composite-trait, the greater the likelihood that natural selection is required to produce it. So when we consider the highly intricate characteristics that make organisms so well-suited to their environments, we may ask how many mutational events were required to create such phenomena. Because, *ceteris paribus*, the greater the number of mutational events that any given complex trait requires, the greater the probability that the mechanism of natural selection was responsible for creating it.

2.5 Lessons from the non-causal and non-creative positions

Let us pause briefly here and recap. In 2.2 and 2.3 we examined the non-causal position: that as causality is exclusively mediated through the individual's proximate interactions with its environment that natural selection's causal powers are illusory. In defence of the traditional view of natural selection as a causal mechanism, we suggested that the apparent problem of introducing natural selection into an already-complete causal account does not arise, and this was not due to any error over the nature of causality, but due to error over mereology. The actions and fates of individuals matter to populations because populations are composed from individuals, the corollary to this being that individuals are parts of populations. We observed that the 'life history' events of individuals given causal exclusivity by the non-causalist position do not cause the population to change in the way that one separate event causes another; the causal histories of the individual 'parts' of a population and the causal history of the population are the same history. The non-causalist preference for 'individual-level' events is guided by explanatory framing rather than causal reality. We further observed, that the error of the non-causal position was in reifying the individual, conceiving of the wider population-level picture as background only, and in considering population-level properties to be merely statistical abstractions lacking in causal efficacy. In contrast the causal view treats the members of a population's life history events even-handedly without giving explanatory priority to any particular individual. Millstein's argument effectively demonstrated the explanatory insufficiency of the individual-level perspective as the population-level effect of trait frequency fluctuation could not be explained through the actions of individuals, but through the causal effectiveness of population-level properties.

Then in 2.4 we examined the non-creative view. Here it was alleged that any causal powers that natural selection does possess are limited to eradicating pre-existing, selectively disadvantageous variation. As with the non-causal argument, we noted that the non-creative argument gives clear priority to events at the level of the individual; although here the relevant properties are the composition and fidelity of an individual's genetic inheritance coupled with the specifics of its ontogenetic development. On the rare occasions that novel characteristics do arise within a population, these will be due to mutation or genetic recombination taking place within, and having effects on, individuals: genetic variation

creates, natural selection differentially propagates. In response, we argued that although natural selection is limited to the distribution of traits and thus to changing the composition of populations, it is through the composition of the populations that the various genetic components of complex traits are brought together and their novel characteristics subsequently produced. As natural selection is a stochastic rather than a deterministic mechanism, its effectiveness in doing this could not be shown with certainty. Our approach was to establish that natural selection substantially raises the probability of the unlikely genetic combinations required for complex adaptations existing in the same population at the same time. Conditional on a probabilistic view of causation, we considered this sufficient to demonstrate natural selection's creativity.

What additional insights can we draw from our combined defences of the traditional, causal and creative view?

Firstly, and most evidently, there is both negative positions' overt focus on the actions and events of the individual at the expense of a wider perspective which recognises the causal relevance of other individuals. The non-causal position denigrates the other members of the extant population that together with the target individual jointly produce population-level effects, similarly, the non-creative position denigrates the target individual's direct ancestors and those individuals that comprised the target individual's ancestral populations. This is more than just a quirk of explanatory framing. If the 'denigrated' individuals in either of these cases had behaved differently, the evolutionary consequences would also be very different. For the causal argument, the relationship is synchronic, the causally-relevant individuals that matter are at a spatial not temporal remove and share with the target individual the property of belonging to a common population. This provides another reason for why the collective behaviour of individual population members cannot be said to 'cause' the behaviour of the population; causes are temporally prior to their effects whereas the behaviour of individuals within the population and the behaviour of the population itself occurs simultaneously. If 'non-target' population members behave differently, then the population behaves differently, so denigrating non-target individuals as explanatory

background rather than recognising them as part of a wider causal mechanism unavoidably, and mistakenly, relegates population-level properties as non-causal and derivative.

For the creative argument, the relationship between the ‘target’ individual - here, specifically, an organism - which genetically inherits and develops, and the other non-target individuals which the non-creative position omits from the allegedly-exclusive site of creativity, are not only the target-individual’s immediate parent/s and direct ancestors, but the other members of the ancestral population, where the variants required to create the target individual’s genetic inheritance were selectively favoured and propagated. Here the relationship is diachronic; the causally-relevant individuals that matter are at a temporal remove, but notice that it is again a ‘population-level’ property that is causally responsible and which the negative-position overlooks. Natural selection does not operate by ‘selecting’ or ‘not selecting’, it operates by selectively *favouring* or *disfavouring* and this requires a population for the traits to be differentially-propagated through. To casually account for the creation of complex adaptations- which require the historical propagation of many such traits- we cannot ignore the effects of the members of the ancestral population who are *not* direct ancestors themselves, as their traits *are* responsible for the population-level properties which enabled those members which are direct ancestors to successfully develop and reproduce. In contrast to the position of the causal argument, here the relationship between the target individual and the wider, ancestral, population *is* causal, not mereological. Descendant individuals are not *part* of their ancestral populations, but the selective dynamics of ancestral populations are *causally* necessary for descendant individuals to possess the genetic repertoire that their complex traits require. This is why the nature of causation itself had to be addressed albeit briefly; the relationship between the inheritance and development of complex traits to the properties of ancestral populations is one of cause and effect.

Our second lesson is more of a confirmation. Although both natural selection’s causal and creative abilities have been defended in the current chapter, only the dynamics highlighted by the creative argument may lead to *adaptation* in the sense that over enough rounds of reproduction the members of later populations could be expected - *ceteris paribus* - to have

accumulated more environmentally appropriate characteristics than the members of their antecedent populations. Millstein's causal argument demonstrated natural selection's causal efficacy, but showed only that population-level properties can be causally responsible for population-level effects. As the selective advantage of each genotype was negatively correlated to its representation in population T, the higher the representation of each particular trait, the lower that trait's selective advantage when compared to its alternative. In the absence of changes to the population's environment or the introduction of new genotypes disturbing the system, the 'evolution' of population T could well be limited to recycling the same d- and h-genotype frequencies. This *is* evolution, and natural selection is the cause, but in a quite a limited sense; producing more of a 'holding-pattern' of genetic representation in the population, than a process that moves the population from one former state into another. The difference between the effects of natural selection as characterised by Millstein's approach and the phenomena that we are looking to natural selection to explain can be brought into sharper focus by considering the following:

If the population with the greater selection differential: P2S, were examined at time t1 at the start of the demonstration and then re-examined at a later time of t2, the data analysed from these 'snapshots' might allow us to place P2S at t1 and P2S at t2 in the correct temporal order with some confidence. This would be possible as the representation of the selectively-beneficial X and Y alleles would be expected to increase over subsequent generations at the expense of the comparatively detrimental x and y alleles; later stages of P2S are likely to have more. If a similar operation were attempted on population T featured in Millstein's causal argument however our confidence in placing the data from the different periods in the correct temporal order would be much-reduced. Although selection changes the relative frequencies of the d and h genotypes in that population, the representation of these genotypes is cyclical, not directional as in P2S. We would not be able to tell - at least using the frequencies of the h-types and d-types alone - if one snapshot of the population developed from or gave rise to the other.

As is well known frequency dependence does not always operate in this way, traits may also exhibit positive frequency dependence when their selection advantage increases the higher

their representation in the population- this would promote the directional selection witnessed in P2S. But the negative frequency dependence described in population T highlights an established scepticism concerning not what phenomena natural selection can explain, but what phenomena we can *expect* natural selection to produce (Parker & Maynard-Smith 1990, Sober 1984, Endler 1986, Gould 1998). Natural selection is clearly in evidence in population T - so much in fact that the example was chosen by Millstein with the explicit aim of demonstrating natural selection's causal potency. But equally clearly, as the changes in genotype frequency described in population T lead only to cyclical evolution, this particular type of natural selection cannot lead to adaptation, and therefore cannot lead to adaptive traits. We rightly expect that organisms undergoing adaptive evolution to be markedly different from their predecessors, and in cases of cyclical frequency-dependence *simpliciter* this will not be the case.

If we wish to accurately capture natural selection's ability to equip organisms with complex adaptive traits as per our stated aims in this chapter then it is clear that the arguments given so far will need considerable supplementation. To be sure, our creative argument demonstrated that the unlikely genetic combinations needed to produce complex adaptive traits require natural selection to promote their retention and propagation through the wider population. Natural selection is clearly necessary for this to happen, but even the *non-cyclical* causal influence of natural selection as demonstrated in our creative argument is not sufficient for this and falls short in providing an adequate explanation for the evolution of complex adaptive traits.

Recall our query towards the end of 2.3. Do we consider an explanation for: i) the fact that populations contain individuals which fit their environments more suitably than other individuals, to be the same explanation for: ii) the fact that individuals have certain features which fit their environments? Our creative argument speaks to the first of these explananda: population P2S features four types of individuals {xy, Xy, xY, XY}, the origin of the X and Y alleles is mutational and therefore accidental, but the proliferation of these alleles through their Xy, xY and XY phenotypes is explained by the selective benefits these traits provided in their given environment. This is why the P2S population features individuals more suited

to their environment $\{Xy, xY, XY\}$ than are other P2S population members $\{xy\}$. But this is a long way from explaining why individual organisms such as cacti, scorpions and birds have the incredibly complex and intricate adaptations that they do.

Although we have shown that natural selection can promote the representation of certain traits within a population, and can understand that the retention of these traits coupled with further mutation and (in sexual species) genetic recombination can lead to more complex traits that without natural selection would be vanishingly unlikely, we have provided no additional qualitative connection between the earlier trait $\{xy\}$ and its 'descendant-traits' $\{Xy, xY, XY\}$ other than the fact that these descendant traits are selectively beneficial. This is not enough. For the creation of complex adaptive traits such as spines and wings, later stages must not only provide a selective advantage *per se* to the stages preceding them, the selective advantages of the later traits must, we suggest, be due to them interacting with their environments in quite a specific way.

Let us unpack the claim. We know that in population P2S, according to fitness scheme FSS, the relative fitness of each trait is: $XY > Xy = xY > xy$. But to know if this dynamic may contribute to the evolution of complex adaptive traits we also need to know the *empirical reasons* for each of these traits having the relative fitnesses that they do. Knowing the comparative selection benefits for each of these traits is not enough. Suppose for example that the population P2S is a population of bacteria, and that the xy genotype provides a level of environmental resilience, perhaps partial immunity to a harmful bacteriophage pathogen. Suppose further that the introduction into the population of the X and Y alleles provides Xy and xY organisms with a heightened immunity to this pathogen, so explaining the fitness advantages of the new alleles through their improved resilience, and additionally that the superior fitness of the XY genotype is granted by virtue of XY bacteria enjoying an even higher immunity to the environmental pathogen. Here we can understand how some of the individuals in the population have features which fit their environment as the selective advantages of each genotype maps onto their empirical consequences in a direct and appreciable way: xy is selectively inferior to Xy and to xY because it provides inferior immunity - xy is more vulnerable and therefore more prone to infection than are both Xy

and xY ; XY is even more resilient than both Xy and xY , and is therefore even less vulnerable to infection by the pathogen than are Xy and xY . We explain the selective benefits of genotypes by their average effects; but what makes the selective benefits here conducive to the adaptation of complex traits is that they coincide with performative increases in a narrowly-construed *function*: here immunity to a specific pathogen.

Clearly, the alignment of selective benefit with ‘improvement’ in any given effect like this is far from guaranteed. The empirical circumstances for the fitness benefits of each genotype may be very different. Given different empirical circumstances: an alternative population $P2S^*$ for example, the xy genotype might have provided organisms with a degree of immunity as xy did in $P2S$, while Xy in contrast might not effect immunity at all in $P2S^*$, but provide instead an increase in fecundity; xY likewise may not produce any differences in immunity in $P2S^*$, but instead increase metabolic efficiency; and XY could have provided no differences in *any* of the these abilities in $P2S^*$ but produced substances harmful to non- XY individuals’ survival and reproduction.

Notice that if this were the case, although the accumulation of genetic variation might follow the same evolutionary pattern as illustrated in $P2S(\{xy\}, \text{then } \{x \text{ mutates into } X, y \text{ mutates into } Y\}, \text{ then } \{\text{both } X \text{ and } Y \text{ existing in the same organism}\})$ and retain the same fitness scheme (FSS: $XY > xY = Xy > xy$), the results would be very different adaptively. The haphazard empirical nature of the fitness benefits of each combination in $P2S^*$ - immunity, fecundity, metabolic efficiency, production of non- XY harmful toxin - impedes the evolution of complex adaptive traits that we set out at the beginning of this chapter to explain. For that to happen, natural selection must not merely promote cumulative adaptation, but promote cumulative adaptation of a specific *sort*.

For the spines of cacti and the stings of scorpions to evolve, the selection pressures need to be such that new variants of the such traits functionally outperform their predecessors in quite a specific way. We can see the beginnings of this in $PS2$ when the genotypes map onto immunity, where the xy - Xy/xY - XY individuals strengthen their immunity through through

intermediate steps, but not in PS2* where the fitnesses advantages of the various genotypes are haphazard and unrelated.

2.6 Chapter recap and conclusion

In 2.1 we set out this chapter's main task: defending the traditional view of natural selection as the causal mechanism for the complex adaptive features of organisms and recognised this as a preparatory task in terms of this dissertation's larger project. In 2.2 we discussed the non-causal position but in 2.3 found its arguments unconvincing; focusing on the undoubtably causal interactions of individuals proved to be an unsuccessful strategy in light of the fact that individuals are parts of populations, not the causes of populations. Casual interactions at the level of individuals *are* causal interactions at the level of populations. In 2.4 we examined the non-creative position but also found it implausible, and for similar reasons. Novel traits arise because of the interactions between genetic inheritance and ontogenetic development- properties of individual organisms - as the non-creative view contends - but the reasons for an individual's genetic inheritance involve the selective dynamics of the populations of its ancestors; again the wider perspective was overlooked in preference of events occurring at the level of the individual. In 2.5 we synthesised our observations from the previous sections and noticed that although the the cumulative retention and propagation of unlikely genetic combinations is a necessary step for the creation of complex adaptive traits - and one that natural selection is responsible for - natural selection *per se* is not sufficient to account for the 'fit' of organisms to the environments that they inhabit. For that we need to more closely examine the empirical reasons for some combinations enjoying a fitness advantage over others.

With natural selection's causal and creative potency secured, our focus in chapter 3 will move to the type of natural selection required to produce the complex traits which motivated our defence of the traditional view at the start of this chapter. This will involve engagement with a well established area in the philosophy of biology: the concept of biological function,

and perhaps surprisingly, also involve engagement with a very well established area in philosophy more generally: how to relate normative values with natural facts.

Chapter 3

Normativity, Function & Progress.

3.1 Introduction

The use of metaphorical language is endemic to evolutionary biology. Behaviours have ‘strategies’, organisms have ‘goals’, phenotypes have ‘designs’ and traits have ‘purposes’; as is well recognised (Francis 2005, Godfrey-Smith 2009, Okasha 2018b) the application of these terms to describe biological phenomena can be far from philosophically innocuous. But the use of intentional concepts like these are not the only type of metaphor that evolutionary biology invites. Adaptations of some species are described as ‘superior’ to those of others; some species and taxa are thought to be more ‘successful’ than their rivals; and despite general acknowledgement that thinking of some species as ‘higher’ than others has little empirical warrant, its use still echos in the contemporary evolutionary literature. The *evaluation* of nature occurs as a matter of course. Some of these evaluative terms are more easily construed naturalistically than others: we might understand the ‘success’ of certain taxa - perhaps mammals and insects for example - as referring to intra-taxon diversity or geographical distribution, or perhaps the taxon’s endurance over evolutionary time. Any normativity in these statements would seem to be of a harmless, cosmetic variety as the empirical basis underpinning them could be restated in purely descriptive terms without loss of understanding. Other types of evaluation are more problematic. Darwin was well-aware of the fallacy in thinking that nature could be ranked according to ‘highness’ for example (Darwin 1872)⁹: when we think of certain species or organisms as ‘superior’ or ‘higher’ the evaluation appears to be of a different and more questionable stripe.

The difficulties of reconciling normative ‘values’ and natural ‘facts’ have an equally distinguished philosophical heritage. New arrivals to this rich seam of analytic philosophy are typically directed towards Hume’s ‘Fact/Value Gap’ (Hume 1739) and Moore’s ‘Open Question Argument’ (Moore 1903), while the more contemporary-orientated reader will

⁹ We return to this problem in Chapter 5.

almost certainly be concerned with attempts at providing ethical attitudes and propositions with a naturalistic explanation (Mackie 1977, Railton 1986, Nichols 2004, Vavova 2015). On the one hand we have matters of empirical fact, such as the roundness of the Earth, the atomic weight of hydrogen, and evolution by natural selection- descriptions pertaining to how the world *is*, while on the other hand there is a quite distinct form of proposition: ‘murder is wrong’, ‘cats are better than dogs’, and ‘help should always be preferentially given to family members’- which concern value judgements directly or directives based upon the desirability of outcomes and how the world *ought* to be. The philosophical difficulty lies in the negotiation of how naturalistic statements concerning facts which - as the majority view has it - should be value-free, relate to normative statements concerning attitudes or directives which are value-laden. Compatibilists have limited choices: either construe normative properties to be natural facts about organisms like ourselves; perhaps a set of rules as to how they must act if they are to promote social cooperation (Sterelny & Fraser 2017), or consider normative properties to be fully ‘real’ but independent of the agents capable of holding them (Boyd 1988).

Our aim in this chapter will involve a reconciliation of sorts between certain natural and normative phenomena, although the ‘direction of travel’ will be the reverse of that usually taken. Typically, analysis commences with normative phenomena: with moral, ethical, perhaps aesthetic attitudes, and attempts to explain these as adaptive in some way so as to provide them with a natural evolutionary origin. Our aim here is different. Instead of trying to ‘naturalise normativity’ in this way, we will show how some natural facts allow - and indeed invite - normative interpretation. Moreover, it will be increasingly clear that this approach is not idiosyncratic and does not run counter to common usage. When we want to understand how certain traits provide a selection advantage over others, the empirical facts require us to interpret their performances evaluatively. In contrast to the ‘naturalising of normativity’ are per the more usual compatibilist approach, we will instead be attempting to ‘normativise nature’. This will not quite be the post-Humean revolution compatibilists have been waiting for, but an elucidation of the reasoning implicit in a great deal of evolutionary explanations. We will further show however that this normativity is not merely an explanatory convenience or introduced artificially; the inherent naturalism *and* normativism

of certain evolutionary phenomena can be shown quite straightforwardly: the functional performances of biological traits can be objectively better or objectively worse than the performances of other traits. We are not projecting our values onto nature or transgressing philosophical orthodoxy in recognising this, but describing a real and causally efficacious aspect of the evolutionary process.

Our findings will lean on two very different philosophical frameworks: Christensen's division of normativity into distinct Evaluative and non-Evaluative conceptual areas (Christensen 2012), and the selected effects (SE) conception of biological function (Wright 1973; Millikan 1984, 1989, 1989b; Neander 1991a, 1991b, 1995a, 1995b; Godfrey-Smith 1993, Kitcher 1993, Allen & Bekoff 1995). Despite keeping to the letter of the SE approach we noticeably deviate from its application in two ways; we will dispute one of its oft-suggested consequences on the one hand, and highlight an often overlooked consequence on the other. The SE conception tells what traits did, but it does not tell us what traits 'ought' to do; there is no normativity of the 'Evaluative' kind - according to Christensen's rubric - to be found in evolutionary phenomena although some SE adherents claim otherwise. We identify that the term 'selected' in selected effects refers to a process, rather than an event; the selection dynamic in the SE conception is not one of the effects of traits being selected or being not-selected, but one of traits being favoured or disfavoured due to functional differences in their performances. Holding the SE conception of function requires an implicit commitment to a type of progressive evolution brought about through improvements in the performances of traits' selected effects.

We begin in **3.2** by identifying different types of normativity and noticing that some of these types can be, and very often are, applied to phenomena quite outside the actions of moral agents. Some performances are objectively better than others in a way quite unlike that of ethical or 'Evaluative' judgment. Then in **3.3** we divert our attention specifically to the performances of biological functions. The 'selected effects' (SE) conception of functions is considered to have introduced normative terms into evolutionary discourse but we see here that this attempt is less successful than is usually thought. In **3.4** however, we show that although the more established route to 'normativising nature' through the SE conception

fails, the SE conception does permit the use of normative language in a milder ‘Non-Evaluative’ form. **3.5** integrates these findings with this dissertation’s overarching subject matter: progress in evolution, and addresses preemptively a possible criticism: how we should consider cases where ‘improvements’ in biological function compromise fitness. In **3.6** our attention turns to scope: we explain that the type of progressive evolution described is far from a minor part of evolution- but is a widespread characteristic, traits by and large provide benefit through improvements in their selected effects. **3.7** revisits this chapter’s findings and places them into our larger project.

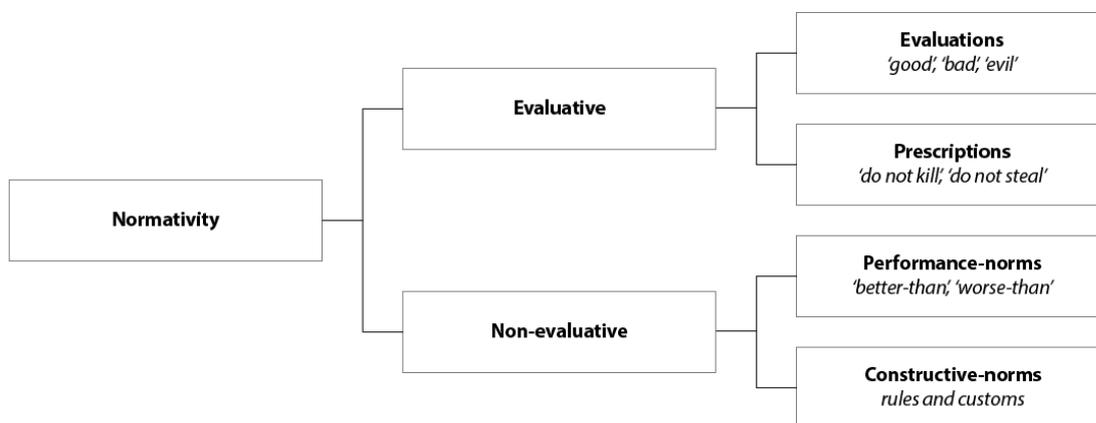
3.2 Naturalising normativity & normativising nature

Attempts to make clear the relationship between normative attitudes and natural facts remain one of philosophy’s staple preoccupations. Approaches to the problem tend to begin with values: an action, attitude or outcome that is considered desirable or undesirable, and attempt to find a natural, non-normative justification for the holding of these beliefs. But the relationship between facts and values can often be equally opaque when considered from the opposite perspective, as certain natural phenomena and evolutionary phenomena particularity seem to invite normative appraisal. This chapter will be orientated with this lesser-addressed mystery in mind; instead of ‘naturalising normativity’ as per the more established route, we will instead be addressing the ‘normativisation of nature’. Identifying certain natural, biological facts and showing that our normative interpretations of them are not mere projections as might be commonly construed, but properties belonging to the phenomena themselves which we accurately describe by performative evaluation.

In preparation for this task we will briefly identify the various species of normativity in order to make clear the nature of our ‘normativisation’ using the broad strokes of Christensen (2012) as a heuristic. According to Christensen we have Evaluative Normativity on the one hand and Non-Evaluative normativity on the other. Evaluative Normativity - concerning values as more commonly construed - covers two sub-categories: *evaluative* terms such as ‘good’, ‘bad’, ‘moral’, ‘evil’, etc and *prescriptions* such as ‘do not kill’, ‘do not

steal’ and so on. This is where philosophical efforts to reconcile natural and normative concepts are typically directed: ‘what natural facts conspire make actions good or bad?’ or dictate that ‘killing is forbidden?’ - projects which attempt to ground the explananda of ethical attitudes and directives within a naturalistic picture of the world, often using social or evolutionary principles as explanans. In contrast Non-Evaluative normativity also covers two sub-categories: what Christensen terms ‘*performance norms*’ such as those which make a participant’s skills and abilities in a given activity poor or proficient, and ‘*constructive norms*’ which specify the necessary rules that must be observed for a given practice or activity to be counted as an instance of a broader type. Clearly Christensen’s demarcation of the normative landscape into these limited categories is somewhat crude but as we will see, it can be put to productive work.

Fig 3.1¹⁰



In contrast to Evaluative normativity, Non-evaluative normativity has been given far less philosophical attention for understandable reasons. There seems little controversy that when we stipulate that moving the king one square at a time is a necessary condition for an activity to count as a game of chess, or that the goal-keeper alone may use their hands to

¹⁰ fig 3.1. A representation of Christensen’s ‘taxonomy’ of normativity. Subsumed under the broad umbrella of ‘Normativity’ itself are ‘Evaluative’ and ‘Non-Evaluative’ strains; under Evaluative normativity lie ‘evaluations’ themselves and ‘prescriptions’; under Non-evaluative normativity lie ‘performance-norms’ and ‘constructive-norms’. (Graphic by Tim Ashton: original to this dissertation)

touch the ball if the game is to be a qualify as game of football - both norms of construction - we are referring to rules that are minus any ethical baggage, more straightforwardly analysed, and of a quite a different stripe to those which prescribe that we should not kill or lie. Performance norms are similarly considered uncontroversial - or at least uncontroversial when applied to *inanimate* processes or artefacts. When a colour television does not produce the colour red, or when there is heavy snow in April, the apparent 'wrongness' of these situations can be readily explained by a failure to achieve a pre-assigned task in the case of the television or by a departure from expectation or regularity in the case of the unseasonal snowfall. So although something of a blunt instrument in the manner wielded here, Christensen's categorisation brings into sharp relief an important distinction which the remainder of this chapter will capitalise on: when televisions malfunction and bad weather strikes any norms transgressed are clearly on the Non-Evaluative, *performative* side of the divide rather than Evaluative. Broken televisions and bad weather can be annoying and unexpected but they are not immoral or unethical.

Matters rapidly become problematic however when we consider deviations from *biological* performance norms: low blood pressure, albinism, hyperplasia for example. The apparent 'wrongness' of these conditions cannot be explained by departure from expectation or regularity so easily. Biological conditions such as these seem to be unlike that of the faulty television as there is no conscious agent in the role of 'assigner' to establish what the biological performative norms are, and so no obvious way that these conditions can be described as not achieving them, but these conditions are also unlike unseasonal weather in that there seems to be something quite distinctive of the compromised performances of biological entities over and above their statistical deviation from what items of a similar type regularly do. Decades of snow in April for example would likely establish this as the 'norm', whereas we would not consider albinism to be any less of a failure to produce pigment if - *ceteris paribus* - cases of it significantly increased in the much same way that we currently disvalue gum disease despite its endemic presence in many species including our own. Regularity or expectation in such cases does not seem to impose biological normality on them. Biological performance norms seem to be different.

Although statistical approaches which do recognise disease as a deviation from species-typical performance retain popularity (Boorse 1975 & 2011, Kingma 2010, Griffiths & Mathewson 2018) our objective here is somewhat tangential to the debate: the compromised performances of both the faulty television - an artefact, and the hypotensive heart- a natural organ, present themselves as examples of the *same kind* of deviation whereas the unseasonal snowfall and the hypotensive heart- both natural events, do not. If we consider the similarity in deviation in both the heart and television to be similar despite the very different natures of entities responsible it would seem that in terms of performance-norms drawing a line between artificial performance-norms and natural performance-norms is drawing the line in the wrong place. When a television cannot show the colour red it fails to meet its performance norm because televisions are *supposed* to display red not because on average televisions do actually show red, when a heart cannot provide adequate pressure for circulation it too fails because it would seem hearts equally are ‘supposed’ to provide adequate pressure for circulation, not because hearts on average actually do provide adequate pressure. It may not be supposed to snow in April either of course, but the use of the term is recognisably distinct here in that the supposition *is* underwritten by regularity to some degree. If it frequently snowed in April, it would be unclear what point someone who claimed that it was not supposed to happen was trying to make.

Although the deviation of the television and the deviation of the heart may strike us impressionistically as the same kind of deviation, our basis for the appraisal of one is quite clearly different from our basis for the appraisal of the other. We know the standard of performance assigned to televisions and other artefacts as *we do the assigning*; if the screen fails to project an expected colour it is failing to perform to the standard for which it was explicitly designed. The performance norms for biological traits are not accessible to us in this way as we had no hand in their production. This problem may strike us as one quite familiar to the philosophy of biology. In much the same way that anatomy requires us to reappraise the concept of ‘design’ minus a designer (Dawkins 1986, Ayala 2004), and adaptive behaviour requires us to re-conceive of ‘reasons’ minus a reasoner (Dennett 1990), the phenomena of biological performance norms would seem to require an ‘assignment’ of appropriate performance minus any assigner. The issue of apparent intention and reasoning

in evolution is typically approached by drawing an analogy between the processes conscious agents undertake when genuinely designing and reasoning, and the processes through which natural selection produces biological phenomena *as though* it could design and reason (Godfrey-Smith 2009, Okasha 2018b).

Nature does not of course design organisms as we do artefacts by the use of models or by stress-testing materials ahead of construction, nor does nature identify a problem such as maximising foraging returns with minimal time and energy expenditure and from this desiderata calculate the optimum strategy. But the lack of actual intention in the evolutionary equivalents of these ‘design’ and ‘reasoning’ processes is often considered relatively unimportant compared to the phenomena produced (Dennett 1995, Dawkins 2005). Although the characterisation of evolution as a problem-meets-solution type of process can be contentious (Gould & Lewontin 1979, Levins & Lewontin 1985) it is clear that questioning the design specifications of the mammalian eye, or asking for the reasons why some birds fake injury in the presence of predators are perfectly warranted empirical enquires *despite* their metaphorical presentation. When we suggest that the ciliary muscle of the eye is designed so as to lengthen or shorten the lens allowing for differential focus, or that the reason birds pretend to be injured is to distract predators away from their vulnerable young, these are suggestions that can be investigated and we may hope definitively answered; despite ciliary muscles not having actually been designed, and birds not actually reasoning about the effectiveness of their distraction techniques. If someone were to claim that asking design questions about the ciliary muscle or reasoning questions about birds’ behaviour were somehow misconceived and unwarranted, they would not be only be denying the legitimacy of the analogies but handicapping biological investigation.

Given the reality of design minus a designer, and of reasons minus a reasoner, how may we go about interpreting the ‘norm’ of performance of any biological trait or behaviour minus an explicit assignation? This may be very obvious in *some* cases: an eye that is severely injured or diseased and cannot focus at all will surely not be performing to its ‘assigned’ standard, nor would the distraction technique of a bird positioned where an approaching predator could not see it. But often things will not be so straightforward. The subtlety of

many traits means that their effectiveness can be measured to a much finer degree than a simple and binary on or off. A bird faking an injury for example may be quite convincing in its behaviour so as to be indistinguishable from one genuinely injured, whereas another may not be quite as talented but still fake injury to a lesser but still appreciable extent. Anatomical traits may perform to different degrees as much as behavioural ones. Even slight differences in phenotype and appearance may translate into large differences in effectiveness. This can be understood particularly clearly where mimicry and camouflage are concerned for example; one mimic may look considerably more similar to its given model than does another, just as one organism may blend into their environment more invisibly than others do. Even to those skeptical that normativity can be anything other than a projection of values onto natural phenomena, there is an unavoidable sense in which a mimic which resembles its model to the point of indiscernibility, or the organism camouflaged so as to be undetectable are performing these tasks *better* than those less-disguised, and that these differences in ability belong to the organisms themselves rather than to an unwarranted normative interpretation.

Notice that the skill of a bird's distraction technique or the effectiveness of a mimic seem to be open to the performatively good or bad interpretation in much the same way as the performances of athletes or gymnasts are, but with a major difference. When we consider the superiority in performance of one athlete over another, or of one gymnast over another, we can be reasonably confident that we are making a factual rather than an impressionistic or subjective judgment because we know what *the rules of these practices are* - in Christensen's parlance, we know their constitutive norms - and so are able to adjudicate their performances within the context of these accordingly. Comparing the performances of the traits of organisms in this way is more problematic for obvious reasons. It is not immediately clear how we can contextualise the performances of the traits of organisms in the same way as we do the performances of sports players as it is unclear what the analogous 'rules of the game' are for the activities of traits and behaviours which we did not invent. The constitutive norms which stipulate that the performances of different players are performances in the same activity it seems are not provided to us by the natural world in the same readily-accessible way.

For example: at the time of writing the current holder for the fastest 100m is Usain Bolt (clocking in at 9.58 seconds, Berlin 2009). Clearly the whole notion of a professional sporting event is somewhat artificial: the distance to be run, the medications competitors are permitted to use, position of the competitors' feet on the starting blocks and so on; there seems to be little 'natural' about such events. But once these constitutive norms have been decided on and enforced, the performances of the competitors are not evaluated but *measured*. We know that Bolt was the *best* sprinter in that particular race, because we know that the rules of the game were enforced, so each competitor's performance in the event was a performance in the same activity, and because there was only one way to be the best sprinter in that race which was to run *faster* than everyone else. 'Better' in this example equates straightforwardly with 'faster'; as time and distance are empirically measurable quantities, so are the relative performances of the competitors. Notice that although there is non-evaluative normativity at work here, the norms are not being employed to judge the relative *performances* of the competitors - empirical measurement is doing that - the norms evoked here are *constitutive* norms employed to ensure that the performances of the competitors are performances at the same identically-constituted activity.

This objectivity in measurement can be carried over to many different types of performances once the constitutive-norms of the activity have been established and enforced. Consider the following: we stipulate that we wish to maximise the storage capacity of a specified area using the minimum amount of materials so as to accommodate a maximum specified load. Various engineers suggest to us various solutions. It seems a misnomer here to describe the process by which we arrive at the best solution as an evaluative 'judgement'; much like the performances of the athletes in the 100m were timed, given the establishment of the competition's constitutive-norms the best solution in these circumstances is one that we can more accurately be said to *calculate*. The design specifications here are of course decided by the preferences of the specifier, but once these norms have been established the superiority or inferiority of differing solutions according to these specifications are no longer a matter of subjective preference.

If we find the use of comparatives in the above engineering and sprinting examples to be problematic it will not be due to any obvious *normative* difficulties. We may have doubts about how similar design competitions and sporting events are to adaptations like mimicry and camouflage given their artificiality, but there is little in our use of comparative terms such as ‘better’ or ‘best’ here that would disturb Hume or Moore. Following Christensen’s Evaluative vs Non-Evaluative distinction it is clear that any normativity concerning the engineering competition and the sporting event falls squarely on the non-evaluative side of the divide. Fulfilling design criteria or setting world records are not in themselves *evaluatively* ‘good’ or ‘bad’ things to do, nor are these *prescriptive* courses of action that engineers or sprinters are ethically obliged to follow; the norms here concern *construction* and *performance*.

Competitions must be played fairly, in that competitors must be competing at the *same thing*, but to ensure this they must also be organised fairly too. Events such as the one won by Bolt are run on the understanding that stopwatches are accurate, that the track is of suitable material and length and so on. These conditions may not immediately strike us as ‘constructive-norms’ in a metaphysical sense, but notice that the organisers of such events need to ensure these and other conditions for the race to qualify *as* ‘the 100m’; results from other ‘100m’ races could not be compared if their track were longer or shorter or of different material, or if their measuring equipment were not suitably comparable. The coherency of there being a ‘100m world record’ depends on this similarity. The engineering commission is similarly constituted. It will be applied for on the tacit assumption that it contain all the relevant information required for its successful completion- if the request for the best use of storage space omitted details such as the area in question being on a submarine, or that the construct had to be made from reclaimed materials this would be a constitute a *different* commission from the one written down. An engineer who lacked this information and as a result lost the commission to a rival applicant possessing it could not be blamed for offering an interior solution; her solution would be designed to solve a different, less-specific project.

In terms of performance the engineering commission and the sporting event are very likely to produce their ‘best’ entrants as they are both contrived to discover these very things. Performance in these cases is underwritten by the very limited number of ways that their results are appropriately measured. These restrictions mean that comparative success is objective and definitive: Bolt (9.58 secs) being *faster* than Gay (9.71 secs) or Powell (9.84 secs) is like Bolt (1.95 metres) being *taller* than Gay (1.8 meters) or Powell (1.9 metres). Given the specifiable norms of construction to guarantee ‘fair play’ there can be no genuine disagreement about who is fastest or tallest because coming first in these cases equates to satisfying a small collection of empirically verifiable facts.¹¹

As can be appreciated from their brief implementation, Christensen’s two Non-Evaluative forms of normativity are fundamentally connected. Constructive-norms are the conditions required for something to be what it is: the rules of the 100m ensure that everyone is running the ‘same race’, the rules of chess ensure that opponents are playing the same game; performance-norms on the other hand concern how well or poorly a game or activity is being played or performed. Notice that we are only able to appraise the playing of a game or the performing of an activity if the constructive-norms for that game or activity have been established; a comparison of performance is only possible when applied to performances of the same thing. Knowing the performance-norms for ‘artificial’ practices such as sprinting or chess-playing is relatively straightforward as the construction-norms for activities were explicitly designed with competition in mind. Accessing the performance-norms for biological traits and behaviours on the other hand would seem to be considerably more problematic as their constructive-norms are not given to us in the same readily accessible way. If we wish to ‘normativise nature’ non-evaluatively by referring to differences in

¹¹ This not the case *per se* in competitive sport and the reasons for this are informative. Being fastest in a sprint, highest in a jump or furthest in a throw are of course very difficult things to achieve, but they are not difficult things to *judge*—as long as ‘judging’ here means recognising the best performance rather than ensuring that the rules have been followed. Anyone competent can judge certain events and arrive at the same results because of the objective nature of the performances. In events where the performances of competitors are *not* objectively measurable: gymnastics, diving, figure skating etc. then judgment *is* required, which is why these competitions have multiple judges in order to ameliorate subjective bias. The difference here can be appreciated by considering the immediate crowd response in track events; unqualified ‘lay’ spectators are able to straightforwardly perceive the win. This is not to denigrate non-track events to any degree, in many respects these sports will be harder as there is a greater range of abilities that participants must master; the point here is that judging these events is much harder too, because translating the diverse abilities it takes to be, say, a winning gymnast into a smaller number of comparative scores is a skilled task in itself.

biological performance, we will first have to ground these activities naturally, to find their construction norms, so as to ensure their comparison in warranted.

Our positive response to this challenge will be the focus of 3.4. But first we will consider whether the grounding of biological performances and activities supports the suggestion that biological traits could be said to have ‘obligations’ - or perhaps less contentiously, whether naturalising biological function warrants the claim that the functions of biological traits are performances or activities that traits are ‘supposed’ to have or are ‘meant’ to do. As we can appreciate through Christensen’s taxonomy: if there are activities that biological traits *ought* to do, the normativity in play here would seem to be of the strong Evaluative variety similar to ethical directives such as you ought not to kill or steal etc, if true, this would clearly be a significant victory for the project of normativising nature. As we will see however, this stronger claim cannot be supported.

3.3 Normativising nature #1: The ‘obligations’ of biological traits.

On the rare occasions when attempts to ‘normativise’ evolutionary phenomena have been explicitly undertaken (Millikan 1984, Papineau 1984) the approach has been to interpret the functions of physiological or anatomical traits as quasi ‘obligations’. Hearts have the function of pumping blood, eyes have the function of seeing, and the spines of cacti have the function of protecting the plants from thirsty herbivores. If we accept that these are accurate statements, then depending on how we understand the term ‘function’ there appears to be only the slightest of terminological shifts between saying that ‘hearts are for pumping blood’ and ‘eyes are for seeing’ to ‘hearts are meant to pump blood’ and ‘eyes are meant to see’. We start with an innocuous statement about what adaptive traits do and very quickly arrive at what looks very much like an explicitly normative directive. As we will see, this is not quite right. Functions *can* be used to demonstrate the inherent normativity involved in some evolutionary phenomena but this is not due to traits somehow ‘dishonouring’ their functional obligations.

How we should understand the phrase ‘the function of the eye is to see’ has largely coalesced around two distinct approaches¹²: the Causal Role or ‘CR’ conception and the Selected Effect or ‘SE’ conception (Wright 1973; Cummins 1975; Millikan 1984, 1989, 1989b; Neander 1991a, 1991b, 1995a, 1995b; Godfrey-Smith 1993, Kitcher 1993, Allen & Bekoff 1995). According to the CR conception of function, the function of a trait, part, or item is the contribution that trait, part or item provides to the functioning of the higher system of which it is a part. Defenders of the CR conception often point to its generality as this definition can apply to either biological or artificial functions: the function of the pixels in your television are to broadcast in a full spectrum of colour as this activity contributes to the functioning of the television as a whole; the function of the eye in those organisms which possess them is its contribution to the organism’s visual system, and ultimately to the fitness of the organism. According to the SE conception by contrast, the function of a biological trait is the effect for which that trait has been naturally selected: the function of the eye is collecting and processing light as this is the effect for which eyes have been historically selected. Due to its reliance on natural selection the SE conception can only apply to explicitly biological and evolved traits- a bullet which its supporters are happy to bite as it also provides a reliable means through which biological traits may possess a function despite not being currently able to perform it. The function of the eye is to collect visual information on the organism’s environment, as that is the effect for which eyes have been historically selected; this remains the case whether an eye is currently performing this function or is not.

Notice that despite these profound differences, biological functions defined by CR and those defined by SE are largely coextensive; they rarely disagree about *what* the function of any particular biological trait is and only then in somewhat esoteric circumstances. Both versions would agree that the function of the heart is to pump blood, that the function of the

¹² Wouters (2003) Identifies and defends four different ways that an enquiry into function might be spelt out in biological study. We might here be asking: ‘what does it do?’, ‘how is it used?’, ‘how is it useful?’ or ‘for what effect was it selected for in the past?’ (Ibid p636). I am sympathetic to Wouter’s analysis, but the distinction between the first and second of these senses, which hinges on what an ‘item’ does by itself *independently* of any selective advantage or role that the item plays in a wider system, need not concern us here. Wouter’s third ‘biological advantage’ construal of function is already- to my mind at least- distinctly normative in that it describes the effects of traits that are ‘useful’ to their owner. This will not concern us here either. We will achieve little by extracting values from a concept which is already evaluative.

eye is to see, and that the functions of leaves are to receive sunlight and exchange gases provided the ‘higher system’ required for CR functions is the fitness of an organism. This alignment mostly holds because when CR and SE advocates do disagree it is typically over *why* traits have the particular functions that they do. For defenders of the CR view, the functions of leaves are collect sunlight and exchange gases and the function of eyes is to see because these are the roles that these traits play to contribute to the higher functioning of the system of which they are a composite part: here the fitnesses of the plants and animals and which possess them. For defenders of the SE view, these traits have they functions they do because ‘seeing’ is the activity that eyes have historically performed and for which they were naturally selected, and leaves have the functions of gas exchange and sunlight collection because these are the activities for which these traits were selected in the ancestors of today’s vascular plants.

It is only when a trait’s current contribution and its past activities differ that these descriptions come apart. If a trait originally adapted to perform function ‘A’ now contributes to the fitness of its owner by producing an effect that is not A, CR will identify this new activity as a different function: ‘B’. SE will not identify this novel effect as a new function until it has been selected for and this requires supplementary conditions: the trait’s new effect must be inherited along with the trait which in turn will require some environmental stability, the new effect must provide descendent organisms with a selective advantage, and selection must be cycled through multiple generations so that an explanation for the trait’s prevalence in the ancestral organism’s decedents will include the beneficial ‘selected effect’. Only once all these conditions have been met will SE and CR reconcile in recognising the novel effect as ‘function B’.

The relevance in how each of these approaches awards traits with their specific functions however has wider consequences than mere disagreement over marginal cases. The CR interpretation of the sentence ‘the function of the eye is to see’ will tell us that the eye is *now* contributing to the fitness of the organism possessing it by seeing- but that is the limit of the information conveyed. We cannot read-off from this CR application a reliable track record of eyes contributing to the fitness of the ancestors of contemporary organisms. What

eyes *used* to do in terms of CR functionality is considered irrelevant, it is what they do now that is important. Notice that this, metaphorically speaking, absolves the eye of any functional responsibilities in that there is nothing that eyes are ‘meant to do’. When the functionality of any given trait is contingent on the fitness benefits of its current effects, there can be no ‘established’ standard to which traits of the same type should be expected to conform. This is often considered a problem for the CR account, as it appears there is no straightforward way that traits can *dysfunction* without first establishing what their correct functioning is (Neander 1995c, Schwartz 2007)

For SE advocates the statement ‘the function of the eye is to see’ contains a much wider implicit claim as function allocations do not directly refer to the effects that current traits have on their owners’ fitness, but to the vast number of similar effects ancestral traits granted to ancestral organisms, specifically those which provided a selective advantage and so aided the traits’ propagation. As successive rounds of historical selective advantage have established what the function of a trait is, unlike with CR functions, *current effects can be compared and contrasted with their previous performances*. This permits a concept of dysfunction which is straightforward and intuitively agreeable: if a trait is currently producing effect A and was selected for producing effect A then the trait is functional; if a trait is currently not producing effect A, but was selected for producing effect A then the trait is dysfunctional¹³. So eyes that are unable to see are dysfunctional; not because they compromise current fitness, but because the selective advantages of eyes came by virtue of them seeing; leaves that cannot exchange oxygen and carbon dioxide are dysfunctional: as exchange is one of the effects that leaves were selected for and so on. But note that although individual ‘token’ traits may qualify as either functional or dysfunctional through their

¹³ Notice that for the SE approach, as functions are the effects for which a trait has been selected, confirming the specific functions of traits will be a historical matter, and in common with all historical events - evolutionary or otherwise - the evidence required for this confirmation may be empirically inaccessible. The weight of evidence required for an adaptive explanation is central to the wider adaptationism debate, and has been given considerable philosophical attention (Gould & Lewontin 1979, Orzack & Sober 1994 & 1996, Brandon & Rausher 1996, Griffiths 1996, Sober & Orzack 2003, Sober 2008, Forber 2010). Orzack & Sober (1994 & 1996) consider an ‘ensemble approach’, where the strength of natural selection in the evolution of traits could be discernible through a comparison of predictive models which allow for natural selection only against the predictions of multiple alternative models whose included evolutionary influences are more varied (although see Brandon & Rausher 1996 for methodological criticism). Sober & Orzack (2003) update this approach by suggesting that the importance of natural selection may be appreciated in descendent traits given background assumptions on their shared ancestry.

comparison to past ‘type effects’ when we label a given effect *as* a ‘function’ we are making an existential claim about a series of similar effects that ancestor-traits of the same the type have successively produced and benefitted from. Trait T cannot have function A, unless there is a sizeable heritage of traits ancestral to T whose ‘A-ing’ provided them with a selective advantage. The history of traits ‘following in the family business’ in this way seems to impose certain expectations. If eyes favoured the selection of ancestral organisms possessing them, then this, at least partly, explains why eyes persisted and proliferated. If the selection advantages granted by the possession of eyes were by virtue of them seeing, then we seem to have possible warrant for claiming that seeing is what eyes are *for*- they are present in today’s organisms *because* they allowed their ancestors to see. The existential claim of a tradition of traits T with the functional effect A appears to justify the prescriptive claim that A-ing is what Ts are *meant* to do.

It is perhaps more striking than is usually recognised that stipulating that functions are the effects for which traits have been selected appears to simultaneously bestow on them qualities long thought exclusive. Functions now appear to be *normative* in that they inform us of the effects that traits are ‘meant’ to have, but equally functions are now *natural* in that they describe real properties of organisms entrenched firmly in the material nexus of cause and effect. However, as we will now see, SE’s acknowledged claims to normativity - at least in this respect - are considerably less secure than is often assumed.

Firstly, normativity. Or at least the *type* of normativity suggested to permeate SE functional statements. If we return to Christensen’s earlier distinction it is clear that it is of a different stripe to the normativity appropriate for the fastest sprinter or the most efficient design solution. If Ts are ‘meant’ to A this does not describe *performative*-normativity, as it does not compare the A-ing of Ts with the A-ing of anything else, it merely states that they should A. Nor is the statement *constitutively*-normative: it does not require Ts to A in order to count as Ts. This point is given particular emphasis by Neander (1991a) who stresses that traits which cannot *perform* their functions must still *have* their functions if we are not to be prevented from discussing dysfunctional traits for fear of logical contradiction. The claim that ‘Ts are meant to A’ is of a much stronger kind of normativity: it is distinctly

prescriptive. SE functions appear to stipulate what traits *ought* to do. Our earlier suggestion that hearts which do not maintain blood pressure have more in common with faulty televisions than they do freak weather events appears partly justified by this development. Televisions *ought* to broadcast in the full colour spectrum because that is what they are for—doing this explains why there are colour televisions; perhaps hearts equally *ought* to provide adequate pressure because that is what they are for—doing this explains why there are hearts. ‘The weather ought not to snow in April’ is not a prescriptive statement of this type as previous snow-free Aprils do not explain why later Aprils largely share this quality.

But closer attention to the prescriptive ‘ought’ at play here shows it to be somewhat mysterious. Consider again the terminological shift from: i) ‘ancestral eyes gave the organisms possessing them a selective advantage by seeing’ to ii) ‘eyes are for seeing’. We can recognise i) as a straightforward SE function claim, past effects (seeing) which provided a selective advantage become the trait’s (eyes’) function. Statement ii) is also a function claim, but unlike i) it is ambivalent about what *sort* of function it considers ‘seeing’ to be. It could refer to a CR function and claim that seeing increases the fitnesses of sighted organisms; it could refer to a SE function, and so be equivalent to i); or it could refer to neither of these, and perhaps cite instead the intentions of a divine creator or similar. Clearly these differences have major consequences for whether we think the shift from i) to ii) is terminologically justified. If the function cited in ii) is a CR function then this in conjunction with i) is an empirical claim: the effects which gave ancestral eyes a selective advantage are the same effects that currently increase sighted organisms’ fitness; if the function in ii) is an SE function then ii) is a straightforward restatement of i); only if the functional ‘is for’ in ii) is genuinely intentional are we straying beyond the remit of describing natural phenomena. So long as we are clear that this is not what the ‘for’ in ii) means, we are either making claim about the similarity of past and present effects or restating the information conveyed in i).

Now consider the shift required to justify this third statement iii) ‘eyes are supposed to see’ or perhaps ‘eyes are meant’ to see’. Unless we have departed from natural explanation and believe the ‘for’ in ii) to be explicitly intentional, the prescriptive qualities of ‘meant’ have

been imported here without warrant. As we have seen, CR functions only describe current effects, and so have no bearing on what traits are ‘meant’ to do *per se*, but neither do SE functions tell us what traits are ‘meant’ to do either. That traits have an ancestry of granting a selective advantage is simply a statement about the effects that those traits had on past organisms, like the lack of April snow it is a *historical* claim. SE functions differ from weather conditions in that past effects of traits account for their current prevalence, but this is a recognisably *causal* claim about the conditional existence of contemporary traits, it is less clear that this is a *prescriptive* claim about what traits are ‘meant’ to be doing. As we can appreciate from Christensen’s taxonomy, the normativity alleged to be in play is of the stronger Evaluative variety, and its application is simply not justified here; at least not without a more explicit explanation of what the ‘meant’ in these statements is intended to convey.

SE advocates might, we suspect, offer the following defence. Artefacts of course *do* have effects expected of them as there are conscious agents like ourselves justifying this by virtue of us doing the expecting; so the ‘meant’ in ‘televisions are meant to broadcast in colour’ does not describe any properties of televisions themselves so much as the properties of those doing the expecting, these could perhaps be characterised as certain mental states. Due to the formal similarities between some artefacts and some traits, it may on occasion be heuristically beneficial to think *as though* traits had ‘obligations’ like it is often thought beneficial to consider anatomical ‘designs’ or the ‘reasons’ for evolved behaviours and adaptations. But notice that if this is the way that SE functions ‘normativise’ natural phenomena, then like the terms themselves the prescriptive aspects of functional effects can only ever be metaphorical. Anatomies have *actual* designs which can be discussed and investigated, behaviours and adaptations have *actual* reasons which we can try and discover; it is quite unclear what investigating the ‘obligations’ of a biological trait would amount to (how obliged the trait is to produce its given effect? The circumstances in which its obligation is lifted?) Once metaphorically prescriptive statements like ‘eyes are supposed to see’ are unpacked the ‘supposed’ disappears and all claims to normativity leave with it revealing a statement about causes and historical events only. If this is right, then clearly the suggestion that SE functions bridge the fact/value gap has been somewhat overblown.

3.4 Normativising nature #2: The performances of biological traits.

There is an alternative means of normativising nature. In common with the failed approach above, the genuinely *normative* aspects of evolution that the SE conception does bring to light are accessed through its ability to establish functions as *natural*. As we saw, SE functions are not relativised to any given system as are CR functions, but are objective features of biological entities which have been caused - by natural selection - and in turn have produced effects - the propagation of the organisms possessing them. This not only permits an entirely naturalistic construal of functions but importantly, also allows us an entirely naturalistic construal of functional performance. To appreciate this we must briefly return to the sports field. Recall that we paused our investigation into *performative* norms by highlighting the objectivity of the competitors' performances. This was demonstrated by the ease with which certain events can be adjudicated. If the output of each performer is a straightforwardly measurable quantity then the comparison of those quantities will itself be a factual matter. What makes some kinds of performances unnameable to direct empirical measurement is the broadness of skills which competing in these events require; there is one way only to win a sprint or a long jump, whereas out-performing your competitors in gymnastics requires speed, strength, dexterity, precision, tempo, and many other abilities as necessary (but by no means sufficient) conditions. There is a clear distinction here between performances that are *measurable* because they are objective, and those which must be *judged* because they require some conscious discernment when comparing performances. We do not need to draw up a comprehensive list, it should hopefully be clear that activities such as sprinting, jumping and lifting are independent of judgement because they are empirically quantifiable. As we appreciated through Christensen's normative taxonomy in 3.2, as long as these activities are regulated and their *constitutive* rules enforced, the better performances in these types of activities will be objective.

With our analysis of SE functions still fresh in our minds the relevance of this analogy should be increasingly clear. As the function of a trait is constituted by its history of

selection we can be assured that the functions of traits are as natural and objective as the traits are themselves. Participants in the functional debate often see this too as a weakness of the CR conception, particularity when applied to biological traits: the CR functions of items appear to be limited only by our imagination in conceiving of certain effects as productive in a higher system (Millikan 1989a). This is emphatically not the case with SE functions; we can be confident that the function of eyes is not to support the eyelid or to itch when tired etc because we can be confident that ‘seeing’ is what eyes were selected for¹⁴. Note that by virtue of telling us what the functions of traits *are*, the SE conception also supports trait-to-trait functional *comparisons* because it informs us that the closer the shared selection history of any token traits, the closer the functions of those traits will be. This bonus ability ‘comes for free’ if we hold that functions are selected effects, as shared etiology ensures that the constitutive-norms normally provided by convention in sporting events have been satisfied.

This claim may initially appear somewhat implausible but consider two full siblings; according to the SE conception of function, the functions of any traits that they have in common will be identical because up until their shared parent(s) produced them the selective history of these traits was identical. This is not coincidental or contingent: for the SE conception shared ancestral selection *necessitates* shared function. Not only do ancestor-descendent traits have functions in common, closely-related traits have functions in common too. Identity of function will hold in this way until the divergent lineages of ‘sibling’ traits undergo enough differential selection so as to make their causal histories distinct. Before a lineage diverges to the extent that they have distinct selection histories and thus possibly distinct functions, related traits will have the same functions, so like the rules which constitute the fair participation in the sporting event, this guarantees that the performances of the traits - if their SE functions have been correctly identified - will be

¹⁴ Identifying the actual SE function of a trait will not always be as straightforward as ‘seeing’ is for the eye and in many cases this procedure may well be beyond our investigative capacities. As Sober (1984) (following Wright (1973)) emphasised, there is an epistemological imbalance here: we may be confident that there has been selection *of* certain traits - as evidenced through their increased populational representation - more than we may be confident of the specific effects of that trait which have been selected *for*. The effect of natural selection may be quite clear, while the precise casual circumstances of the selection remain opaque.

comparable in an objective and principled way¹⁵. Notice that the constructive norms characterised by Christensen as mediated through cultural practice or custom are here established quite naturally. We know Bolt was the fastest, because the rules of the 100m ensured that the performances of competitors were objectively comparable; here we know the performances of traits are objectively comparable as shared causal history ensures that their performances will be performances of the same type.

So if the functional performances of biological traits can be shown to be measurable- like running or weight-lifting etc- as opposed to judged- like gymnastics or ice skating- then we have an alternative, solid basis for the application of normative terms to biological function, and one that does not rely, like the more established normative route, on equivocal terminology. Although the SE view cannot tell us what traits are ‘meant’ to do in Christensen’s Evaluative-prescriptive sense, it does tell us that traits with the same selection histories have the *same function* - even if the specific details of this history are unavailable. So when we do have grounds for knowing what the particular function or functions of a trait are, we will also have grounds for knowing how *well* any particular trait is functioning, and therefore also have grounds for knowing when any particular trait is functioning *better* in comparison to the performances of other traits of the same type.

This can be shown this quite clearly as follows: take two traits T and R belonging to two closely-related organisms (we may for simplicity consider these as allelic variants). As the selection histories of both T and R are near-identical, we know that the functions of both T and R must be the same function. Even if we do not know the empirical facts, we can be sure of this because the etiological perspective granted by the SE conception of function tells us in general terms what the function of the trait is: the selected effect which gave the shared ancestors of T and R a reproductive advantage. We will call this - perhaps undefined

¹⁵ When the precise nature of a trait’s selected effects are beyond empirical reach, this will clearly also place any comparison of trait-to-trait functional performance out of reach also. This epistemological difficulty does not threaten our central point: for there to be a difference in selection benefit, the effects of the traits must be different, and the cause of this difference will either be due to differences in functional performance or due to one or both traits producing a novel effect- even if we are unable to discern which one of these distinct eventualities has occurred. I am indebted to one of my examiners for pushing me for clarification on this point.

- selected effect s . Now consider how we would explain differences in T and R's fitness. We will of course be able to discern if there are differences in T and R's fitness due to their systematic differential reproduction; here we will say T has this advantage. As we know organisms with T have a selective advantage over those with R we know that these traits are having different effects - obviously. But we also know, due to their shared causal history, that T and R have the *same* function. Notice that this would not be possible from the CR perspective. There can be no guarantee for T and R having the same CR function because their contributions to the fitnesses of the organisms possessing them are different contributions. In CR terms this could be understood as T and R performing different functions, whereas the SE conception prevents this: etiology dictates function, and the etiology of T and R is near-identical.

So far we know that both T and R are doing the same thing: s , but that T is doing s differently to R to the extent that it gives T a fitness advantage. Now it may very often be the case that the nature of T and R's function s will remain mysterious; perhaps because their population is hard to access in some way, or if the performances of T and R are quite subtle. But even absent this information we know that s -ing was selected for, and that T is s -ing *better* than R because the effect which defines both T and R as functional is here making T *more* functional. This particular shift in terminology may initially strike us as being equally as incongruous as the previous introduction of 'meant' and 'ought' to impose expectation, but here it is quite warranted. Due to the specifics of the SE account the 'function-of' *equates* with those effects for which the trait was 'selected-for', if T and R have the same selection history, they will have the same function; if T enjoys a fitness advantage over R, then this will be due to T interacting with its environment in a way which is novel or idiosyncratic or it will be due to T performing its SE function, but performing it to higher standard than does R.

This dynamic can be shown more clearly by applying it to a solid example. Suppose we identify T and R as the wing patterns of a population of butterflies, and identify T and R's shared function s as camouflage. We notice that the T-pattern's representation in the population is increasing compared to the R-pattern and conclude that the T-pattern gives the

organisms possessing it a selection advantage. How do we explain this? Given access to the requisite empirical evidence, we deduce from their shared histories that the function of both T- and R- patterns is camouflage, given this we then know that T is doing this more effectively than does R: we know that Ts are *better* camouflaged than are Rs. Note that the 'better' functioning of T here is a straightforward description of an empirical phenomenon in the same way that describing Bolt as a better runner was a description of an empirical phenomenon. It is just the means by which the competitions are constituted that is distinct: comparative performance in sporting events is constituted by agreed-upon rules; comparative performance for trait variants is constituted by their shared selection history. Once the competitions have been constituted- then the performances within the competitions are amenable to objective measurement.

This application of 'better' functioning can be applied to evolutionary dynamics quite generally if the functional identity of the competing traits is ensured by their shared history of selection in this way. The function *s* here could just as easily refer to 'anaerobic respiration' and T and R as two closely-related protozoans. We notice the systematic reproductive advantage T-organisms have relative to R-organisms, and - provided the empirical facts on the ground are available - are able to deduce that T and R's function *s* is anaerobic respiration by their selection history. Given that both Ts and Rs were selected because of their benefits to anaerobic respiration, if they are functional now they are still doing this, and if Ts have a selective benefit over Rs, it is because they are functioning better than Rs. Selection for anaerobic respiration was responsible for the spread of the traits ancestral to T and R, and now selection is responsible for the greater propagation of T traits in contrast to the spread of Rs. Unless T is producing an entirely novel effect, this is because the way in which Ts are *s*-ing is superior to the way in which Rs are *s*-ing.

But will this always be the case? And even when it is, will it be as clear as our examples so far have suggested? Someone resistant to the idea that functional improvement may be straightforwardly read-off from a trait's current selective advantage, coupled with evidence of the trait's selection history, might respond with the following observation. All we actually know is that Ts have a reproductive advantage to Rs, and what effects the traits

ancestral to T and R used to produce which gave them a selective advantage historically¹⁶; this does not tell us that Ts are now doing the same thing as Rs, as the effects Ts are now producing may be entirely novel and unrelated to *s*. Perhaps it is not that Ts are doing the same thing as Rs are doing but doing it ‘better’ which provides them with a selection advantage, but that Ts are doing something else *e* that Rs do not do at all. This is a valid criticism, but it only limits the *scope* for functional improvement, it is not our intention here to stipulate that all natural selection of closely-related traits occurs through an improvement in their previously selected effects alone.

This can again be made clearer using a well-known example, the dark-winged moths in the north of England that possessed a selective advantage in the areas affected by industrial pollution. Due to the sudden change in their environment, the darker members of a mixed population of light and dark moths gained an unexpected selection advantage due to the novel camouflaging effects that they, but not their light coloured conspecifics, enjoyed in their newly-blackened environment. Note that the effect of the darker trait here is novel and fortuitous; it is not an improvement of a function that all moths previously possessed and which the darker variant now performs to a superior level in comparison to the performance of the lighter variant, but a change in environment which has suddenly made the darker variant selectively beneficial. Other than the fitness of the darker moths nothing is ‘better’ than it was before the environment darkened. The language of improvement seems misapplied here though the reasoning quite accurate: on many occasions a change in

¹⁶ The issue of which exact period of trait’s selection history should be considered as authoritative for an accurate functional diagnosis remains contentious (Millikan 1989b, Kitcher 1993, Godfrey-Smith 1994) and is thought by some to fatally undermine an etiological approach (Cooper 2002 & 2005). Evolutionary histories, by necessity, cover vast spans of time and what was a trait’s function in the distant past might well not be that trait’s function in its contemporary setting; feathers being a notable example: whatever their original function - now thought to be insulation, or perhaps display - it could not have been flight as the first feathered animals lacked the supplementary anatomy to achieve this. Clearly as our attentions are concerned with the functional effectiveness of contemporary traits, our interests will lie with a ‘modern history’ approach like that of Godfrey-Smith (1994) who identifies the relevant period as the ‘recent past’. Although Cooper (2002) is right to suggest that the selection histories of traits will very often be difficult, if not impossible, to uncover, this is not as damaging to our account as one might suspect. Firstly: by stipulating the recent past as the period in question the loss of evidence will - *ceteris paribus* - be ameliorated to some degree, certainly in comparison to potential evidence located in the ancient past. Secondly: as demonstrated by our leaving the precise selected effect of *s* originally undefined, even if the selection history, and thus functional identity, of competing traits is inaccessible the lacuna is epistemological, the metaphysical point still stands.

environment or appearance of a novel variant will grant a fortuitous fitness benefit for idiosyncratic reasons.

But the arrival of a new *variant* will be accompanied by the arrival of a new *function* only in a minority of cases; environments change and novel variants arise, but we would be overpopulating the category of ‘function’ quite dramatically if we were to think that all environmental or phenotypic novelty produced a new function *Sui generis* each time one of these provided a selection advantage. In the greater number of cases selection will be due to traits functioning to a higher standard than to traits performing functions that other traits do not have at all. As we saw, our exclusively T or R moth population has two distinct variants only, one of which has a function and other of which does not- camouflage is either ‘on’ for the moths of this population or it is ‘off’, but in the majority of cases the composition of a population, and the functional effectiveness of the traits within a population will be considerably more nuanced and varied.

For example: let us now suppose that our moth population is more diverse, and instead of two types of trait it now features three: the dark trait T, the light trait R, and an intermediate grey trait G. We notice that G provides a selection advantage over R, but not to the same extent as T does. In these circumstances analysing the effects of these traits as either ‘functional’ or ‘non-functional’ will be inadequate as we have three values that require explanation. The natural way to approach this would be from the basis that whatever selected effect T produces, G produces less, and R produces even less of - if at all. This might be quite clear in some examples like the trivariate moths where the intermediate grey phenotype is understood as providing intermediate camouflage, but very often trait variance will not be discrete in this way and be manifested on a continuous scale.

Let us now further suppose that this is the case with yet another moth population; one where the appearance of its members now ranges continuously - and we may add, more realistically - across the grey scale from dark to light in the wings of the moths comprising the population; we notice that the darker the trait the higher its fitness advantage. Coupled with environmental information, the colour of the foliage, the visual capacities of the moths’

predators and so, we would be quite justified in explaining the proliferation of darker variants by reference to them being *better*-camouflaged; to explain the compromised fitness of lighter moths due to them being *poorly*-camouflaged; and to explain the fitnesses of those traits of intermediate shades as due to them being *better*-camouflaged than some, but not as *well*-camouflaged as others.

Notice that unlike the use of ‘highness’ to describe anatomical or behavioural complexity, or the use of ‘success’ to describe a taxa’s endurance, diversity or distribution, the use of the performative terms ‘better’ ‘poorly’ and ‘well’ here is not metaphorical but accurately describes how these traits’ differences in appearance affect the fates of organisms possessing them. Camouflage is a real effect, it can be empirically quantified, and moreover the qualitative variation in this ability has casual efficacy, the population’s environment is such that the more effective a moth’s camouflage, the higher its chances of survival and reproduction. Natural selection is here promoting *better*-camouflaged moths at the expense of those which are concealed less effectively, so over time - *ceteris paribus* - we would expect those with the darker traits to proliferate as darker traits provide better camouflage than do the traits with lighter colouration¹⁷. If will still feel that normative language has somehow been smuggled into this description of empirical reality without warrant, consider how we might explain the differential fitness of the wing colours in this population *without* the use of ‘better’ or ‘worse’ or their performative equivalents. The explanation - if possible - would strike us as peculiar and contrived. When we wish to understand not only the diversity of trait variants but how their diversity can translate into differential fitness, the relevant effects will be those that have historically promoted their bearer’s survival and reproduction: the trait’s functions. When we describe the various differences in functioning which different traits produce, ‘better-than’ and ‘worse-than’ are not normative terms which

¹⁷ Our expectations may of course be confounded as there are many reasons why the improvement in biological functioning described here could be halted or reversed - even discounting the non-selective evolutionary influences of mutation, migration and drift. The developmental basis for the darker wings for example may have adverse effects elsewhere in the organism; selection for the darker phenotype perhaps ameliorated by selection against other pleiotropically-linked traits. Traits do not reproduce by themselves independently, they come as a package deal together with other traits of the organisms possessing them. If darker moths are slower moths or less fecund ones due to developmental inevitability, natural selection could eliminate the cryptic trait as a consequence.

are imported without license, they perform an essential role in explaining the nature of the evolution which has taken place.

3.5 Functional improvement and progressive evolution.

Let us pause briefly to review this chapter's findings so far. Through the application of the SE conception of biological function we know that the effects for which traits were selected - their functions - are real, these are not metaphorical, nor they are not merely projected onto biological phenomena for instrumental reasons. Although this is often seen - in Christensen's terminology - as warranting an Evaluative-normative approach to biological function in that traits have effects that they are 'meant' or 'obliged' to produce, we remained sceptical. Claims of historical selective advantage are simply historical claims, they are not normative ones. Knowing the selected effects of traits however does allow us compare the performances of those traits. As etiology tells us what the function of a trait is, it also tells us that traits with the same etiology have the same function. This provides the constructive norms required for the comparisons of biological performances to be justified as performances as the same activity. When traits with the same function are producing different effects, one or both are dysfunctional or both are functioning but to different degrees; when this is the case one will be functioning better than the other. So applying the SE conception of function in this way not only provides us with sufficient grounds to establish functions as natural and to recognise the same function as possessed by instances of different trait tokens, it also provides us with a natural source of normativity - at least of the performative kind. History tells us what functions of traits are - thus satisfying the constitutive norm required for appropriate comparison - and given access to this history we can know when traits are performing their functions to a better or worse degree.

So far, those sceptical of this chapter's initial claims concerning the normativisation of nature might remain unimpressed. We started by demonstrating that the performances of athletes and designers are 'better' than others - a performative-norm -, and only then because the competition - maintained by constructive-norms - is highly artificial. Following

this, although we managed to secure the performances of biological traits conceptually using their histories of selected effects as natural constructive-norms, the most that this shows is that the comparisons of performances of biological traits or behaviours are similar to the comparisons of athletic performances in being ‘better’ or ‘worse’ and common sense, plausibly, does this for us already. Notice however that the ‘thin’ type of non-evaluative normativity provided by the naturalisation of functional performance has significant consequences for this dissertation’s overarching project.

As we saw, differences in functional effectiveness can have significant evolutionary effects, the prevalence of darker wings in our moth population with continuous variation being a paradigm example. If our reasoning up to this point is sound then it would seem that in cases like this, where the evolution of the population is driven through the functional superiority of certain traits in comparison to performances of others, that the evolution in question could be rightly considered as ‘progressive’. Members of the population, by virtue of their better-functioning traits gain a selective advantage due to those traits, and *ceteris paribus*, given repeated rounds of this dynamic the functional performance of the traits under selection should increase. We have evolution - as the composition of the population is changing, and we have progressive evolution - as the differential selection of certain traits is driven by their functional superiority. Natural selection is here influencing the population to *improve* in that function. In our idealised moth population, as camouflage is empirically quantifiable, whether one member of the population’s camouflage is better than another member’s given their shared environment will be a factual enquiry, and one that given adequate evidence empirical analysis could provide an answer for.

Notice that conceiving of progressive evolution in this way avoids two acknowledged pitfalls. Firstly, it is not impressionistic, relying on an abstract notion of ‘highness’ or ‘superiority’ or similarly-vague judgments concerning which parts of the natural world present themselves as more advanced than others. And secondly, although this type of progressive evolution is *produced* by natural selection, the progress here does not *consist* in making individual organisms or species simply ‘fitter’ - as some naive readings of

evolutionary progress might contend - but in improving their performance at a real and potentially quantifiable functional ability.

The profound difference between progressive increases in functional improvement and increases in fitness *per se* can be brought into sharp relief by considering the following: depending on the overall size of the population, the fitness of the average moth after repeated rounds of selection for darker variants might be higher, the same, or perhaps even lower than that of the ancestral moth population featuring continuous variation. Over time the population would get *darker*, but the average moth would not necessarily get *fitter*. The selective benefits of darker traits are contingent on the population's dynamics. The darker traits possess a selective advantage over the lighter traits, but as we observed in chapter 2, traits are not simply selected or not-selected, they are selectively-favoured or disfavoured; as the darker traits saturate the population their advantage over the population-average pattern peters out as the average itself becomes darker. Being selectively-favoured requires the presence of alternative traits less-likely to promote survival and reproduction in the population, without this variation natural selection has nothing to select from.

So after repeated rounds of natural selection the average moth could be no fitter than the average moth in the ancestral population, nor could the darker traits be of any selective benefit, so in many ways the evolution which we are here describing as 'progressive' leaves the population very much as it was. The average moth will be *darker* however, and the crucial point here is that provided the environment remains constant the average moth in the later population will be *better camouflaged* than the average moth in the ancestral population. Progressive evolution - or at least, the type of progressive evolution that we are defending here - does not make individuals 'higher' nor does it necessarily make individuals 'fitter', but it does make individuals better at certain functions than their ancestors were. It is difficult to see how this chapter's observations so far might be considered controversial. It is certainly unclear what a description of evolution making species 'higher' might mean, and partially unclear what would be meant if we were to claim that evolution made species permanently 'fitter', but we have been quite clear about what 'better' means, and evolution certainly can make populations 'better' at performing certain functions. The later moth

population possesses superior camouflage to that of its predecessors, and evolution by natural selection is the mechanism responsible.

So although the type of normativity provided by the SE conception of biological function fails at providing Evaluative normativity, and the non-Evaluative normativity which it does provide may seem somewhat pedestrian, it does however provide a solid conceptual basis for thinking that natural selection can in certain circumstances be progressive. Not all natural selection will be progressive in this way, but when the traits of organisms provide a selection advantage by virtue of them performing their selected effects at higher standard these traits will be *improvements* on those which preceded them, and the evolution responsible will be *progressive* as a result.

We will finish this section by heading-off some potential problems.

Firstly, and for the sake of clarity, it may be instructive here to make a specific point about what we are *not* claiming. Our positive claim is that in the circumstances described, where traits with the same function have different functional performances, one trait will be performing better than the other, and that when this dynamic is sustained the evolution produced will be progressive. We will return to the prevalence of this type of evolution shortly, but note that we are in no way condoning run-away adaptationism in claiming this¹⁸; it would be perfectly consistent for us to regard natural selection as a relatively unimportant factor in evolution for example, and concurrently hold that when natural selection *does* occur through functional improvement in the way described that evolutionary progress occurs.

Secondly: so far in our examples functional improvement and selective advantage have coincided, but the relationship between the two might not always be so straightforward.

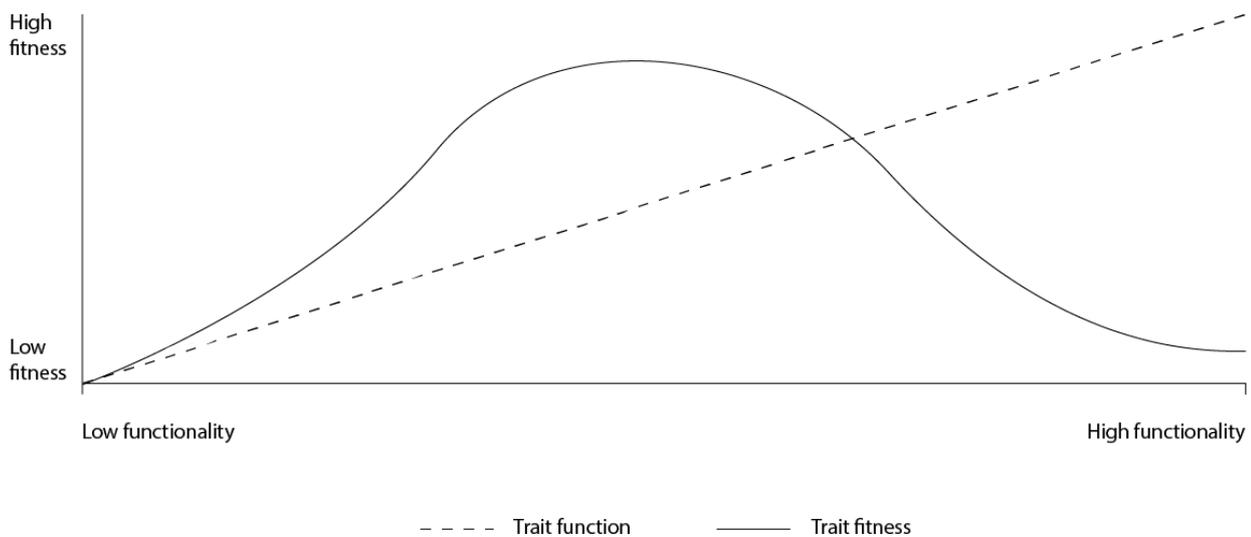
¹⁸ ‘empirical’ adaptationism to be precise (Godfrey-Smith 2001). The other two varieties of adaptationism addressed in the wider debate: ‘epistemological’ and ‘methodological’ are more concerned with the status of natural selection as a guiding investigative or organising principle in evolutionary science, than with the prevalence of natural selection as an evolutionary cause. Although supporters of one type of adaptationism generally give the others sympathetic treatment.

How might we deal with situations where the two were in conflict? As the darker variants spread through the moth population their selective advantage is ‘in step’ with their improvement in functional effectiveness, but the simplicity of this example might mislead us into thinking that increases in a given ability will always be selectively-beneficial. (It is difficult for example to imagine a situation where the moths in question might become *too* camouflaged.) But consider what would happen in alternative circumstances where increases in functional performance result in the ‘improved’ traits no longer providing a selection advantage and start to become selectively *disadvantageous*. This time instead of the population being comprised of moths, our test population will be comprised of birds; specifically nocturnal birds who rely on their sensitive hearing to hunt and navigate in low levels of natural light. We know that according empirical data, appraised through the SE conception of function, that the functions of the birds’ ears is to detect vibrations through the air - sound - as these were the effects for which these traits were selected in the birds’ ancestors. Greater sensitivity to sound and a greater capacity to discriminate between different sounds may mean both a functional improvement in the trait’s performance and provide those individuals possessing the augmented trait with a selection advantage over those individuals in the population which do not; but this increase in sensitivity will only be beneficial up to a certain point. To the extent that the additional sensitivity aids in hunting, predator avoidance, navigation etc it would provide a fitness advantage, but ever-heightened sensitivity to sound might rapidly *lower* the traits’ fitness, perhaps through overwhelming or distracting the individuals possessing it. (See *fig 3.2*)

Fig 3.2¹⁹

Should this not unrealistic situation occur, then according to our understanding individual birds would appear to have functionally ‘better’ hearing - but their fitness pays the price. This seems to provide us with a very counter-intuitive result. When evolution is conceived of as a progressive process - however vaguely - one property all might agree on is that the

¹⁹ *Fig 3.2* is a graphical representation of the relationship between the sensitivity of the nocturnal birds’ hearing and its consequences on the birds’ fitness. Birds with little or no sensitivity to sound have a very low fitness, as their hearing improves so does their fitness; functional improvement and fitness here are ‘in step’ - at least initially. As the birds’ sensitivity to sound increases further however, the trait responsible becomes selectively disfavoured as the birds become overwhelmed or distracted by sensory input. (Graphic by Tim Ashton, original to this dissertation).



fitness of organisms undergoing the evolution should increase or at least remain unchanged at the time the evolution is taking place; it should certainly not be lowered as it could well be in the birds with over-sensitive hearing. Whatever we think progressive evolution is it should not make things *worse*.

Keeping our conception of evolutionary progress as functional improvement intact, there are three ways to approach this problem. Approach a) would be to ‘bite the bullet’ and hold that functional improvement ‘trumps’ decreasing fitness, and that those birds whose hearing is sensitive beyond the point where it provides a selective advantage are still ‘better’ at this particular function than their less-sensitive conspecifics, and claim that as far as progressive-evolution as functional-improvement goes this is what counts. An unattractive option given that it runs counter to how we might think of progress in virtually any other context²⁰. Approach b) would be to claim that once the trait’s sensitivity compromises fitness, rather than increasing it the trait ceases to be ‘functional’ and becomes dysfunctional. As we are here adhering to the SE conception of function, by considering the function of a trait to be those effects which have provided a historical selective advantage,

²⁰ The incongruity here between evolutionary ‘advance’ and selective benefit is somewhat reminiscent of the late 19th century claims that the large size of the antlers of stags were due to progressive and directional ‘orthogenetic’ forces, promoting growth and acting in opposition to natural selection by making the stags’ antlers heavier and parts of their wooded environment less accessible thus compromising their fitness (Ruse 1996). For those holding this view natural selection was considered a secondary mechanism, if not dismissed altogether.

we might plausibly claim that if the trait's current effects no longer provide the selection advantages of its predecessors then the trait is no longer performing its function as evidenced by its history of selection. This approach may seem more appealing than approach a) but notice that b) is precisely what we cannot claim. If the functions of traits are stipulated to be the *precise* effects for which they were selected, then *any* deviation from these effects could be considered dysfunctional; not only those increases in ear function which cause disorientation in our bird population, but also the deepening in coloration which provided the darker moths with better camouflage. It seems that in both the moth and the bird populations that the *effectiveness* of the traits' functions has increased, but that the *type* of function of both the blackened wing-patterns' and the augmented ears' has remained the same. If we deny functionality in one of these cases consistency would demand that we deny it in the other.

This leaves us with approach c): maintain that both the birds' over-sensitive ears and the moths' darkened wings are performing the effects for which they were selected - and so are functional - but that in the case of the moths where functional improvement is 'in step' with selective advantage the evolution is progressive, whereas in the case of the birds where improvement in function and selective advantage come apart the evolution is not progressive. This approach may strike us something of an ad hoc amendment, but notice that once the functional performance of the over-sensitive trait starts to compromise the fitness of the individuals possessing it, the reasons for the trait's maintenance or further propagation in the population change significantly. If the over-sensitive trait were maintained or continued to increase in the population we would hold 'evolution' *per se* responsible, but as natural selection should eliminate the over-sensitive trait let alone promote it, the nature of the processes causing this unlikely event would be arbitrary and erratic. Our reluctance in regarding improvements like those which produced over-sensitive hearing as 'progressive' is surely warranted, but perhaps this reluctance is motivated more by the random and arbitrary nature of the non-selective causes required for the trait's continual persistence than to the fact that the trait itself is selectively disadvantageous. Our motivation for thinking of over-functional traits as non-progressive may not due to their effects -

although undesirable - but due to the erratic nature of the causes required to maintain and propagate them.

Let us unpack this claim. Consider how we would explain the persistence of ‘over-functional’ traits like the birds’ hyper-sensitive hearing in a population featuring other traits which are less-sensitive and thus more selectively-beneficial. We know that natural selection cannot be responsible for the maintenance or propagation of the over-sensitive trait as it decreases the fitness of the individuals possessing it, so if we wish to provide an explanation for this in evolutionary terms our options are limited. The ‘over-functional’ trait’s representation in the population could be due to the fortuitous differential survival and reproduction of those individuals which possess the trait relative to those which do not; it could be due to the genetic basis for the trait persisting in the population through repeated mutation or genetic recombination; or it could be due to individuals without the over-functional trait leaving the population or individuals with the over-functional trait arriving in the population. In short: drift, mutation, recombination, or migration. Out of these possible causes - or combinations thereof - only recombination could be construed as entirely non-random in respect of the over-sensitive trait’s representation. By contrast, as we saw in chapter 2, while the mechanism of natural selection acts probabilistically, it does not act randomly in this way. When increases in functionality are selected for we have a non-random, *mechanistic* explanation for the augmented traits’ propagation, but when traits’ functional improvement becomes disadvantageous, we must rely on non-mechanistic, random - or principally random in the case of genetic recombination - causes for the traits continued representation should it persist or increase. So it is not the counter-intuitive effect brought about by ‘progressive’ improvements causing a reduction in fitness which should lead us to dismiss their evolution as non-progressive; but because the persistence of these traits in the absence of positive selection would be due to random and arbitrary causes.

With our thoughts returning to the issue of ‘drive’ identified in our introductory chapter as a recurrent property in previous accounts of evolutionary progress; whatever we think progressive evolution is, there appears to be wide agreement that it should not come about *accidentally*. We address this issue in greater depth in chapter 4.

3.6 The prevalence of functional improvement

As we saw above, a trait's improvement in functional performance and a trait's selection advantage are distinct and can come apart. Improvements in function can be accompanied by decreases in fitness - as evidenced in our example of the nocturnal birds' over-sensitive hearing - and increases in fitness can be accompanied without any concurrent increase in an established function - as evidenced by our example of darker moths gaining, rather than improving in, the function of camouflage due to environmental effects. When functional improvements compromise fitness, their evolution ceases to be progressive due to the change in the nature of the processes responsible for the 'over-functional' trait's persistence - from probabilistic to near-random. When a trait's functional performance and a trait's selective advantage do coincide however we have a strong grounds for claiming that the evolution of the trait is progressive. In such cases natural selection is making the trait better at its given function.

Those *still* uncomfortable with our evaluative approach to evolutionary biology in this way might at this point consider a strategy of containment. Although functions are real - they might admit - and traits' improvements in their functional performances are real, and that when this effect is produced non-accidentally we have warrant for thinking of the evolution responsible as 'progressive' this is a very minor part of evolution and our cherry-picked examples are unrepresentative and tailored so as to produce a progressive result. Functional improvement may be real, but it is rare, and it should not be used as a perspective through which to view the effects of evolution generally or even the effects of natural selection *per se*. So with the legitimacy of progressive evolution secure philosophically we will turn to the question of its prevalence, and whether progress in evolution should rightly be considered as commonplace or exceptional.

Firstly we can note that although the main result of natural selection might well be the elimination of disadvantageous variation as referred to in chapter 2, populations featuring

continuous variation like our final ‘grey-scaled’ moth population are far from an aberration. Populations frequently feature substantial diversity in trait-variation, and while a percentage of this may have little or no effect on fitness, when trait differences do have a differential effect on fitness we will understand this as being caused in a limited number of ways. Traits with higher fitness will either be producing an effect which is new and unrelated to their selective history, or performing the effect for which it was selected but in a manner that is more selectively-beneficial than do other traits of the same type. As we saw, selectively-beneficial effects may arrive *de novo* but this is surely the more-exceptional explanation for the selective advantage of a trait-variant than a seldom-broken rule. It is not unreasonable to suggest that when traits feature variation in form that the reasons for their differential fitness benefits could be due to these traits possessing variation in functional effectiveness, rather than due to their effects being unique or idiosyncratic.

Consider again the function of the heart: to pump blood around the body. No-one disputes this, and certainly not adherents to the selected-effects conception of biological function such as ourselves. But when it is claimed that hearts were ‘selected’ because they pump blood around the bodies of the organisms possessing them, a naive interpretation could understand this as describing a historical population similar to our original moth population which featured either light or dark variants exclusively and where the dark variants were camouflaged and the lighter variants were not; only here it is organisms possessing a heart gaining a selection advantage over those members of the population lacking one. This is clearly incorrect. Hearts do not arrive in the population *de novo* granting the organisms which possess them a selection advantage over those members of the population which do not possess them; if other members of your population have hearts and you don’t you would not make it to term. Clearly, when it is claimed that the ‘function of hearts is to pump blood’, this is meant somewhat elliptically. Rather than suggesting that the presence of hearts *simpliciter* provides a selection advantage over the absence of hearts, what is fully meant by this statement is that certain hearts - or evolutionary precursors to hearts - possessed a selective advantage to other hearts because they pumped blood around the body in a *certain way* while the other hearts did not. It is not simply the possession of a trait and the possession of the function produced by that trait which provides a selection advantage, it

is the possession of a trait which functions in such a way so as to promote the fitness of the organisms possessing it over those members of the population whose traits do not produce the selectively-beneficial effect to the same beneficial *degree*.

The differential functional effectiveness of the pumping of the heart and camouflage of the moths' wings are by no means idiosyncratic examples as can be appreciated in the following less metazoan-centric case. The function of the flagella found in some bacteria is primarily locomotive; the structure either flexes or rotates and the organism is propelled forward as a result. The selected effects conception allows us to be confident that this is the function of the flagella if it can be shown that this effect was selectively beneficial over the bacteria's evolutionary history. Clearly the claim that the flagellum was 'selected' does not refer to a time when there was a bacterial population featuring members with flagella in direct competition with members lacking flagella and the flagella-possessing members won out. As we appreciated in chapter 2 variation in a population is ultimately sourced from mutation with genetic recombination as a secondary mechanism, but the random nature of the variation produced by these processes at any one time would almost certainly not produce a complex trait such as a flagellum from scratch. Complex traits such as this do not simply arrive fully formed, they evolve incrementally and gradually over repeated rounds of random variation and natural selection. As the evolution of the flagella was incremental, we cannot explain its selective benefits by a straightforward comparison of the fitness of bacteria which possess the trait in comparison to the fitness of bacteria that do not. This would be more of a creative exercise than an empirical one; involving the calculation of the selective dynamics of an ancestral population which did not exist.

So the function of the flagella is locomotion due to its history in providing a selective benefit by virtue of this effect, but the bacteria with flagella were not selectively-favoured over bacteria absent the trait. If the dynamic in the ancestral population was not flagella-powered locomotion vs the absence of flagella-powered locomotion, what was it? Clearly, the variation in the ancestral population was more subtle; some bacteria had flagella - or precursors to flagella - which provided those organisms with *improved* locomotive abilities in comparison to the reduced abilities granted by alternative flagella present in the

contemporaneous population. Selection did not favour flagella over non-flagella, just as selection did not favour hearts over non-hearts or - in the moth population with continuous variation - camouflage over non-camouflage; selection favoured flagella *better* able to perform their function over other flagella, hearts which provided *superior* circulation, and wing-patterns which provided *greater* crypticity.

If our reasoning in this chapter is correct, it would seem that progress in evolution is neither a baseless metaphor nor a normative 'curio' in the otherwise naturalistic practice of evolutionary biology, but is characteristic of the processes required to create the complex adaptive traits endemic in the living world. Not all evolution will be progressive in this way: as not all evolution is driven by improvements in traits' selected effects, and not all natural selection will be progressive in this way: as not all natural selection will be driven by improvements in traits' selected effects either- as the sudden arrival of camouflage in our original bivariate moth population demonstrated, traits may gain a selection advantage by producing novel effects rather than improving on those which they already possess. But when the selection advantage of traits is due to their functional superiority, the evolution of that trait will be progressive; the members of populations undergoing progressive evolution of this kind will better at their selected effects than were its pre-evolved ancestors.²¹

3.7 Chapter recap and conclusions

We begun this chapter by noting that the prevalence of metaphorical language in evolution reaches further than the characterisation of biological phenomena as intentional; terms more familiar to normative discourse are often employed to describe certain species, organisms or their traits. Sometimes the use of metaphor can be directly replaced by empirical

²¹ This contrast between traits performing the effects for which they were historically selected, and performing alternative effects which now promote fitness fortuitously is not an original observation. This topic was approached explicitly by Gould & Vrba (Gould & Vrba 1982) in their distinction between adaptations (traits which are currently performing the effects for which they were originally selected) and exaptations (traits which currently produce a fitness-enhancing effect for which they were not originally selected) and was recognisable in our earlier discussion of the differences between the SE conception of function which identifies functions which their historically selected effects and the CR conception of function which ignores history in preference to the contribution a trait plays the functioning of a higher system: in biology, typically the fitness of organism possessing it.

description, although its use on other occasions remains problematic: traits do not have ‘obligations’, traits do however have functions- their selected effects. As the SE conception of function secures the identity of a trait’s function naturalistically it is possible to tell when traits are performing their functions. From this it is possible to tell if the functioning of traits is comparatively better or worse; and from this we may discern when evolution is progressive. This is not a rare event; to the extent that natural selection operates through differences in the functional performances of traits evolution will be progressive.

Our confidence that we were applying performance-norms appropriately was secured by the ability of the SE conception to provide the required constructive-norms. We argued that the functions of closely related traits will be the same function, as functions are selected effects and the effects for which closely-related traits were selected will be the same, we may not know what the functions of many traits are - the details of selection being out of empirical reach - but when we do have access to this information we are warranted in their comparison. Our findings in this chapter relied heavily on both Christensen’s normative distinction, and the SE conception of function but despite keeping to the letter of the SE approach, we disputed one of its oft-suggested consequences on the one hand, and highlighted an often overlooked consequence on the other. The SE conception identifies past selected effects and thus establishes a ‘constructive norm’ - as described by Christensen - concerning their functional identity through doing so, but it does not tell us what traits ‘ought’ to do; there is no normativity of the ‘Evaluative’ kind - also according to Christensen’s rubric - to be found here, although some SE adherents claim otherwise. On the other hand the term ‘selected’ in selected effects refers to a process, rather than an event; the selection dynamic in the SE conception is not one of the effects of traits being selected or being not-selected *simpliciter*, but one of traits being favoured or disfavoured due to functional differences in their performances. Holding the SE conception of function requires an implicit commitment to the this type progressive evolution brought about through improvements in the performances of traits’ selected effects.

With some of the possible consequences of natural selection addressed in this chapter, chapter 2’s findings can be viewed from a fresh perspective. There we defended the

traditional view of natural selection as a process which is both causal and creative, and found the non-casual argument to be unfounded and based upon a preference in explanatory framing, rather than a difference in casual reality. Individuals interact with their environments, and *ipso facto* populations of individuals interact with their environments. The contention of the non-creative view was also discovered to be unfounded: natural selection is not merely eliminative, but is instrumental in maintaining the representation of favourable traits in the population, and through this promotes the combination of favourable traits that without natural selection would be exceedingly unlikely. Now through this chapter's findings we can appreciate how the theoretical perspectives of natural selection in chapter 2 might play out empirically.

Consider the reasons for, and the results of, natural selection's retention of favourable traits and its contribution in their otherwise unlikely combination. Recall in chapter 2 we observed that the arrival of the XY combination trait was made far more likely due to the selective advantage of X promoting its retention and propagation prior to the subsequent arrival of, and combination with, Y. In light of this chapter's findings, let us label the beneficial effect which X provides to the individuals possessing it as v . How might we explain any selectively-beneficial effects brought about by the combination XY? Through the perspective of function provided by the SE conception, we would either characterise these effects as novel if the XY combination were doing something new and different to v , or we would characterise these effects as the same effect v as provided by X alone. If the effects of the XY combination were not characterised as v this would be an example of the *creativity* of natural selection, but it would not be an example of *progressive* evolution as although XY brings something new to the table, it does not improve what is already there. Whereas if XY provides a selective benefit over X, and does this by virtue of the same type of effect v , but by producing that effect to a different standard, then the natural selection would still be creative - it has still combined X and Y after all - but in these circumstances the evolution would be progressive. The combined trait XY is producing the same effect v to an improved standard, and natural selection is the cause.

Our grey-scaled moth population provided a brief example of what the results of successive rounds of this type of progressive evolution might produce. The traits responsible for the incrementally-improving selected effects should - *ceteris paribus* - proliferate through the population, improving successive populations' functional performances. This was easy to appreciate in the idealised moth population as better-functioning mapped directly onto darker colouration, but in other circumstances the relationship between improved functioning and appearance or anatomy may be less straightforward. Consider traits which require a much more diverse range of changes in appearance and anatomy than simple colouration for an increase in their functional output to be achieved. In our nocturnal bird population for example, to the extent that their heightened sensitivity to sound provided a selection advantage we recognised its incremental increase to be progressive; but the anatomical changes required to produce heightened functioning in the birds' auditory systems will be far more complex and intricate than mere differences in colouration. This is clearly the case in our example of the better functioning heart also; an improvement in the selected effects of an organ as anatomically complex as a heart will require far more 'refined' changes than the cosmetic ones required for the wings' increased camouflage ability.

So natural selection can be creative, and natural selection can be progressive, and some progressive increases in functional ability will require more complicated types of variation in the traits which produce them than do others. As we concluded in chapter 2, the greater the number of unlikely combinations required for any given trait, the greater the probability that natural selection would be required to produce it. In the light of this chapter's observations we can appreciate that complex traits such as ears and hearts, gills and spines which provide organisms with a greater adaptive 'fit' to their environments require not just natural selection *simpliciter* to explain their unlikely production, but natural selection of a special sort. As we saw, the nature of the selection history for hearts, flagella and other complex traits is not one of such traits providing a selection advantage over their absence, but of them providing a selection advantage due to their functional superiority over other rival traits whose functioning is comparatively inferior. The selection advantages responsible for the evolution of complex traits such as these will be progressive, in that the

functional identity of the trait undergoing selection is maintained but its functional performance is *improved* upon, and the process must be maintained over substantial amounts of evolutionary time. We could imagine the variation required for the moths darkened wing colouration to be achieved reasonably quickly, the variation required to produce complex traits such as auditory systems and hearts would be far more considerable, and require a much greater span of progressive evolution as a result.

We will return to progressive evolution over longer timescales like these in chapter 5: how evolution within lineages is represented, when progressive representations are warranted, and when progressive representations are in danger of misrepresenting the type of evolution the lineage has undergone. Then discuss more thoroughly the relationship between progress, adaptation and the ‘design-like’ nature of traits in chapter 6. Before this however, our task in chapter 4 will be a slight detour thematically. In the following chapter our investigation is of patterns, their general abstract features and the kinds of processes and mechanisms responsible for producing them. We saw in 3.6 that once traits become ‘over-functional’ and compromise fitness, the accidental nature of the processes responsible for their maintenance and propagation mean the evolution responsible is no longer progressive. Our work in chapter 4 will build on this observation by investigating further the differences in the nature of the processes responsible for genuine patterns, such as the progressive evolution driven by natural selection, and patterns which may appear to be genuine but due to the accidental nature of their production are not. We then apply these findings to some suggested patterns thought to be characteristic of evolutionary history.

Chapter 4

Natural patterns: authentic & quasi

4.1 Introduction

Originally coined in 1953 by the psychiatrist Klaus Conrad, ‘apophenia’ is now more commonly known as the pervasive tendency to seek and identify patterns in random and unconnected information than for its original intended use of characterising the early symptoms of schizophrenia (Brugger 2001 p13). Far from being a diagnostic indicator for mental illness, human sensitivity to recognise patterns has been re-evaluated as an evolutionary adaptation in working order. In a dangerous and uncertain environment, it may be preferable to have your pattern-recognition set to ‘paranoid’ and be disturbed needlessly, than to risk not being alerted due to an insensitive recognition threshold.

Fortunately, contemporary cases of apophenia typically have less life-threatening consequences. They seem to be particularly evident in matters supernatural; the patterns in the stars which motivate astrological predictions, the bet placed according to your lucky number and so on. But not every case is so clearly erroneous and the history of science features a litany of notable examples. Nicolaas Hartsoeker (1656-1725) for example, inventor of the screw-barrel microscope and far from a ‘crank’, believed miniature human-like forms could be seen contained within sperm at magnification; and the astronomer Giovanni Schiaparelli (1835-1910) infamously witnessed the patterns of networked canals irrigating the surface of Mars. Both patterns were widely believed to have their basis in genuine phenomena for some time. As we know, Hartsoeker and Schiaparelli were both later proved to be mistaken when advances in optics showed homunculi and canali respectively to be mirages rather than actual structures, but conversely the history of science also features equally high profile cases where initially impressionistic patterns were later found to be fully justified. The patterns of inheritance discovered by Gregor Mendel (1822-1884) for example were subsequently found to be real effects produced by the genetic configuration of sexually reproducing species, as was the actuality of continental drift as

suggested by the ‘fit’ of the eastern coast of South America into the west coast of Africa. In both of these cases, what started out as perceived regularities were corroborated with incremental but cumulative evidence before eventual vindication for them occurred through the discovery of the underlying mechanisms responsible. Once it was known *how* these patterns were produced, it was known that they *were* produced- by something other than over-sensitive observers.

This chapter will concern itself with patterns, their general, abstract features and the kinds of processes and mechanisms that produce them. This will, at start, be a preparatory exercise. Later we will be looking at some patterns suggested to have been manifest through the course of evolution in order to see how they withstand ‘stress-testing’ by our newly-clarified terms and concepts.

Life on Earth looks very different today than it did at its inception around 3.5 billion years ago. But what exactly has changed? And how should these changes be understood philosophically? As they can be interpreted quite liberally, attempts to address these questions have come from a variety of quarters: evolutionary biology (Huxley 1942; Dawkins 1996a & 1997, Maynard-Smith & Szathmary 1997, Wilson 2001), palaeontology (Simpson 1974, Raup & Jablonski 1986, Fischer 1986, Gould 1998, Conway-Morris 2003) and philosophy (McShea 1998, Shanahan 2004, Ruse 2009, McShea & Brandon 2010, McShea & Simpson 2011). That there *are* patterns of some sort to be found in evolutionary history is generally accepted- and acceptable- but claims concerning the causes of these patterns and their wider interpretation remain contentious. Darwin’s demystification of the teleological aspects of evolution as a kind of ‘force’ driving species, clades or even individual organisms in the case of Lamarck, ever closer to perfection seems to be accompanied by implicit theoretical corollaries. If the mainspring of evolution is the adaptation of species to the shifting goal of capricious environmental challenges, later species should not be qualitatively distinct from those that have gone before; aside perhaps from being better-adapted to the specific environments which they inhabit although even this is far from guaranteed (Van Valen 1977, Sober 1993, McShea 1998). Parochialism would seem to imply that any patterns discernible *in* the history of evolution at scales larger

than local adaptation will not be due to any ‘pattern-making processes’ of the evolutionary theory itself. Such regularities, it is often suggested, must be explained through the actions of non-selective influences such as diffusion away from a minimum level of phenotypic complexity (Gould 1996, McShea 1998), or the requirement that later organisms will by necessity inherit and thus incorporate the developmental systems of their predecessors (Wimsatt 2001).

But the adoption of these corollaries is not merely a matter of philosophical disposition and has been repeatedly transgressed. E.O. Wilson for example claims:

“During the past billion years, animals as a whole evolved upwards in body size, feeding and defensive techniques, brain and behavioural complexity, social organization, and precision of environmental control- in each case farther the non-living state than their simpler antecedents did. More precisely, the overall averages of these traits and their upper extremes went up. Progress, then, is a property of the evolution of life as a whole by almost any conceivable intuitive standard, including the acquisition of goals and intentions in the behaviour of animals.” (2001, p175)

Richard Dawkins, in a similar vein, writes:

“.....adaptive evolution is not just incidentally progressive, it is deeply died-in-the-wool, indispensably, progressive. It is fundamentally necessary that it should be progressive if Darwinian natural selection is to perform the explanatory role in our world view that we require of it.....I do insist that in a majority of evolutionary lineages there will be progressive evolution towards *something*. It won’t however, be the *same* thing in different lineages.” (1997, p1017 & p1018)

Are these authors and others like them suffering from apophenia? And like Hartsoeker and Schiaparelli erroneously identifying patterns that do not exist? Or are they more like Mendel and the early proponents of continental drift in that the patterns they describe are authentic but perhaps currently lack a satisfactory explanation? Or do their claims fall somewhere in

the middle? Perhaps as side-effects of genuine processes, rather than being characteristic of them?

In section 4.2 we will be looking at the general properties of patterns and identifying the different types that our observations of nature present to us. These may be features belonging entirely to the natural world: ‘authentic-patterns’, or may be partially constructed by human observers: ‘quasi-patterns’. The role a ‘mechanism’ plays in discriminating between these different types will be explained, and we will recognise a five-part heuristic which can indicate which sort of pattern we may be dealing with. In 4.3 we will be looking at two natural real-world patterns in order to demonstrate the practical application of these heuristics. Although both of these patterns are extremely familiar, only one of them is authentic according to the procedures spelt out in 4.2, while the other is quasi-. In 4.4 we address the nature of the patterns suggested by Wilson and Dawkins to be present in the history of life. Here again through the application of our heuristic we observe that one of these suggested patterns is authentic, while the other is not. In 4.5 we discuss the wider ramifications of finding what appears to be an authentic pattern in ‘large-scale’ evolution to actually be quasi-, and suggest the properties that a genuinely authentic pattern at taxonomic scales at genera and above would have to possess. Section 4.6 recaps this chapter’s findings and places them in this dissertation’s wider perspective in preparation for the work ahead in chapter 5.

4.2 Patterns: elements, relations & mechanisms

The term ‘pattern’ is an ambivalent one covering a wide range of disparate cases. At one extreme we have patterns like the constellation of Orion created by our association of stars that are actually distantly remote, while at the other extreme we have patterns that reflect empirical reality more accurately: the fluctuations of temperature according to season, the specific proportions of elements in chemical compounds, and the near-universal presence of DNA in living organisms being obvious examples. The impressions and regularities that we call ‘patterns’ clearly form quite a diverse and loosely-defined set. The task of identifying

the sort of pattern we are dealing with is made considerably easier when the mechanism responsible for it is understood to some degree. We can explain the pattern of seasonal temperature changes at various locations for example as primarily due to the orientation and position of the Earth relative to Sun during its yearly orbit. We know that chemical elements combine in predictable quantities due to the specifics of their atomic and molecular configuration, and that the ubiquity of DNA in living organisms is explained by distant but common ancestry. There is no equivalent explanation of this kind for the pattern of stars that comprise Orion's Belt. Of course, we may safely assume its arrangement is due to interactive astrophysical processes that we have some hope of understanding (if not actually calculating due to the vast amount of data it would involve) but an explanation of astrophysical processes will be of a very different stripe to an explanation for why those stars should have settled specifically into a formation *such that* it appears belt-like when viewed from the perspective of a terrestrial observer. The stars themselves are real and their positions are real, and the processes that created those stars and arranged them into their current positions are real, but our interpretation of these elements and their relations as a pattern forming a 'belt' does not reflect empirical reality. It is imposed on them by human observers.

We can observe a three-way split in these types of pattern. The patterns are either: i) entirely imaginary: as in the cases of Hartsoeker and Schiaparelli, where closer inspection demonstrated the observed phenomena to have no basis in reality; ii) partially real: as in the case of Orion's Belt, where what we might call the identified pattern's 'elements' are objectively real, and their configuration is real but where their resulting pattern is artificial and imposed; or iii) real: as in the cases of Mendelian inheritance and continental drift, where both the pattern's elements and arrangement have no dependence on human recognition. Type i) patterns may be of psychological interest to practitioners such as Conrad, but will be of little further particular importance to us. The natures of type ii) and type iii) patterns and the reasons for the differences between them are of considerable importance however.

We are of course not the first to address patterns philosophically. Dennett's highly influential 'Real Patterns' (1991) being by far the most seminal example whose insight continues to motivate discussion - even in areas not originally intended. (Tolliver 1994, Fodor 1997, Wilkerson 1997, Seager 2000, Ross 2004, Joslin 2006, Ladyman & Ross 2007, Chalmers 2012, Sterelny & Griffiths 2012, Burnston 2017).

Although I am unsure as to how a detailed comparison of this chapter's findings might compliment - or indeed contrast - with those of the Real Patterns debate I mention it for two reasons. The first is terminological: however distinct the patterns which concerned Dennett may be from those which will concern us, and however appropriate the label 'real patterns' may be for some of the phenomena we analyse below, the term is too firmly associated with Dennett's philosophy and its subsequent literature for semantic reassignment. So for the sake of clarity and in an attempt to establish our reasoning in what follows as distinct from the Real Patterns discussion alternatives will be used. The second reason is that at least one aspect of Dennett's approach is informative by comparison. For Dennett the nature of the patterns which originally concerned him - certain mental states, scientific abstractions, obscured diagrams - hinge on their interpretation²², he is not concerned, at least explicitly, in associating their ontological status to the means of their generation²³. (This is particularly evident in one of his main examples: a gradually obscured pattern which degrades from five distinct squares into what appears to be random visual 'noise'. Despite each member of this series being produced by the *same* computerised process, with the various levels of noise added artificially, their 'status' as patterns remains mixed.) In contrast, the patterns of concern to us are ones that are easily-perceivable and well-known (and so, quite 'real' according to Dennett's approach), but whose means of generation play a vital role in how we think about them. At best it would be of tangential relevance to Dennett if his five-

²² Although, of course, not *only* their interpretation. Dennett deftly avoids the charge of instrumentalism through his insistence of a 'compressibility requirement', in short: a pattern is real iff there is a more efficient way of describing it (1991, p32). The possibility of compression without loss of information is decided by a mathematical operation, not an interpretive one (Ross, Ladyman & Collier p205, in Ladyman & Ross 2007).

²³ Dennett has consistently tried to avoid any specifically ontological commitments that might be implicit in his work on Real Patterns (uncharitably labelling ontology: '...the caboose, not the engine...' (Dennett in Ross, Brook & Thompson 2000, p358)). Although he has perhaps not been as successful in this as he might have hoped (See Ross' chapter on 'Rainforest Realism' in Ross, Brook & Thompson 2000, and Ross, Ladyman & Collier on Dennett's ontological difficulties in Ladyman & Ross et al. 2007, p196)

square sequence were generated by the hand of an artist or by waves washing stones randomly up on a beach, whereas for us this would make all the difference.

At their most abstract, patterns are regularities, and regularities must be regularities ‘in’ or ‘of’ something. The Fibonacci sequence for example is a pattern made of numbers, the pattern of concentric rings from a stone dropped in water is made from the wave motion on the surface of the liquid, your daily schedule will be a pattern of activities such as working, shopping, eating, maybe childcare and hopefully plenty of sleeping and some leisure activities. Patterns have *elements*, the various things that together form the pattern; the elements of a chessboard are black squares and white squares, the elements of the Fibonacci sequence are numbers, the elements of the wave pattern are the movements of the water.

But elements of a pattern must be related in a certain way. The *relation* between elements may be quite loose and informal such as in the relation between the activities that comprise your daily schedule, alternatively the relation may be quite strict and linear as in the Fibonacci sequence, where the next number in the sequence is the sum of the two preceding. Some other relations may be quite closely linked, but not be ‘sequential’ in this way. This will very often be the case when the elements of a pattern may appear independent, but are in fact jointly produced. We may think of the various clues that together indicate the identity of the murderer: the fingerprints on the gun, the lie told to the police, the secret inheritance etc as forming a very different type of pattern from that of the Fibonacci sequence or your routine. Sometimes patterns may remain undiscovered due to the obscurity of the relations between their elements, like an unlucky detective failing to properly piece together the evidence.

For our purposes the relations which hold between the elements of a pattern must not be generated *accidentally*²⁴. However strongly they impress themselves upon us, there must be a cause or reason to explain why a pattern’s elements are related in the way that they are.

²⁴ As we saw in our chapter 3 example of the nocturnal birds’ over-functional hearing, knowing *how* an evolutionary sequence is generated makes a significant difference in how we interpret the results. When increases in the bird’s abilities were driven by their selection benefits - ie non-accidentally - we thought of the evolution as progressive, when further increases were disadvantageous and not driven by selection, the accidental nature of the non-selective processes responsible for the continuing increase: drift, mutation, migration etc stopped ‘progress’ in its tracks.

We generally regard non-accidental patterns differently from accidental ones even if the patterns themselves are phenomenally indistinguishable - and for good reason. Suppose we used a random-number generator to produce the sequence 1, 1, 2, 3, 5, 8, 13, while this would be just as much an instance of the Fibonacci sequence as it would be if generated by the formula as stated above, knowing the specifics of its generation - the precise means taken to generate each element in the series - would not expand our knowledge in any way, or if it did, it would not tell us anything more about the *Fibonacci sequence*. If we came across the same sequence in some other form however, perhaps in the dimensions of a spiralling shell or whorls of a hurricane, it would be a pattern whose regularity we would certainly want to investigate further as it might tell us more about how the surface phenomena was generated, and perhaps a lot more. How patterns come about can be important because if they are accidental they do not *signify* anything, if they are non-accidental however, they might.

It might be tempting here to place some of the explanatory burden concerning the generation of patterns by stipulating that a 'genuine' pattern be the result of a genuine 'process', and indeed much of the literature concerning evolutionary patterns employs the term in a confirmatory way (Raup & Jablonski eds. 1986; Emerson 2002; Wake, Wake & Specht 2011). We will avoid this. The difference here may seem cosmetic, but as our aim is to emphasise the differences between accidental and non-accidental patterns and as both of these will be the result of processes *per se* the use of the term 'process' will not make the required distinction. There is instead however a subtle difference between 'process' and 'mechanism' that can be capitalised upon. Consider again the stars that make up Orion's Belt, or the numbers of the randomly-generated Fibonacci sequence. It would be wrong, I think, to say that the elements of these patterns were not produced by genuine processes. The stars that comprise Orion's Belt for example were produced by physical interactions that we know, more-or-less, and have a fairly good hope of fully understanding. The elements of the randomly-generated Fibonacci sequence would also be the result of processes so described. I do not know how [random.org](https://www.random.org)²⁵ or other random-number

²⁵ <https://www.random.org>

generators produce the numbers they do, but we know there will be *some* process that allows them to do this reliably and promptly upon request. If the elements of these two accidental patterns are the result of processes, then the patterns produced from these elements are also the result of processes. We might like to think that the three stars that comprise Orion's belt are the result of three *separate* processes (they are quite far apart after all), and that the seven numbers generated randomly are the result of seven different processes, this, at least, is how I would characterise it. But any potential confusion as to what counts as a single 'process' or plural 'processes', or any potential confusion between the kind of process that produced the Fibonacci sequence through application of a formula, meteorological dynamics, or engineering principles, and the kind of process that generated the Fibonacci sequence by accident is ameliorated by the introduction of a generative mechanism.

As with patterns, mechanisms come in many different stripes, and as with patterns we are far from the first to give their application philosophical interest (Machamer, Darden & Craver 2000, Glennan 2002, Woodward 2002, Skipper & Millstein 2005, Barros 2008, Matthewson & Calcott 2011, Moss 2012, Nicholson 2012, Cartwright 2017). Debate continues, although there is broad agreement that all mechanisms are mechanisms *of* some phenomena. They may *produce* phenomena, when a causal sequence produces a distinct end-product, such as when a virus produces the symptoms of a disease, or when an act of parliament produces a piece of legislation; they may *underly* phenomena, such as the mechanisms of the engine that underly an aeroplane's movement; or they may *maintain* a phenomenon, such as the mechanism in your thermostat that keeps the room temperate- a state of affairs that is preserved by the thermostat's mechanism.

Unlike the description of processes, mechanistic explanations separate the accidental from the non-accidental by drawing our attention to the importance of *reliability*. Mechanisms are 'structured' in a way that processes, largely, are not, something which greatly promotes our ability to understand them. A mechanism that maintains the HbS gene in human populations exposed to malaria for example is more reliable and repeatable than a process that produces the same effect. This does not necessary mean that processes are always idiosyncratic

events, or that mechanisms are always repeated. We may easily conceive of a mechanism of which there was only one (perhaps a prototype) and which was used only once or perhaps not at all, and the mechanisms of history are no less mechanisms for being singular.

Discussion of unique or erstwhile mechanisms will inevitably draw our minds towards evolutionary history. Certain seagrasses have a unique approach to pollination (Pettitt, McConchie, Ducker & Knox, 1980) and certain species of frog are able, uniquely, to survive the partial freezing of their body fluids (Layne & Lee Jr. 1995). It is highly likely that these adaptations were produced only once through the mechanism of natural selection- although additionally natural selection may maintain these traits in the relevant populations. That these adaptations are unique does not then demote natural selection to a mere ‘process’. Mechanisms do not need to be *repeated*, but they do need to be *repeatable*, and this requires structure and regularity in the sense that given the same or very similar initial conditions the mechanism *will* behave in the same or very similar ways, not that it has actually done so. The structure and regularity of mechanisms can be understood counter-factually: had there been no natural selection, then these adaptations would not have existed (Woodward 2002), although consistent with our defence of natural selection’s creativity (in chapter 2), mechanisms can also be understood stochastically, as raising the probability of possible but exceedingly unlikely phenomena occurring (Barros 2008). Either way, mechanisms do not need to have their regularity and structure explicitly demonstrated by repeated instances (Glennan 2010).

Additionally we should not conflate mechanisms themselves with mechanistic explanations (Matthewson & Calcott 2017), so some caution is required. Although this practice may be largely innocuous, it obscures the fact that comparatively simple mechanistic explanations can be applied to complex and disparate phenomena, very different from mechanisms as more traditionally conceived. The mechanisms that produce, underly or maintain patterns come in many different varieties, as do the target phenomena themselves. The mechanism that we understand to produce a thunderstorm for example will be very different in nature to the mechanism inside a smoke detector, although both of these are composed from a number of sub-systems or ‘components’ acting in a fairly prescribed manner to jointly create the target phenomena. We can see the components that jointly comprise the mechanism of the

storm are not physically- distinguishable separate ‘parts’ whose interactions are straightforwardly-localised, although a mechanistic explanation for the storm’s behaviour might treat it as such. There is no specific location within the storm’s mechanism where the component of ‘low air pressure’ interacts with ‘high air pressure’ for example, this happens throughout, so in contrast to its mechanistic description the number and type of actual interactions between the storm’s components will be vast and hugely complex. This is not case with a more traditional mechanisms- a carbon monoxide detector for example. Its specific components- the chemical detector, LED, alarm etc are quite discrete, localised and interact in a very limited number of ways. The components of the detector are of course themselves comprised from other parts, and ultimately from the interactions of its constituent atoms, this is a major advantage of identifying the processes responsible for a pattern as a mechanism. Mechanisms ‘*screen-off*’ what can be described as the “concatenation of lower-level causes” (Fischer 1986); if there are processes and sub-mechanisms with their own dynamics operating within the larger mechanism, mechanistic explanations often provide us with the option of ignoring them.²⁶

Mechanistic descriptions provide a natural cut-off below which an explanation need not go in order to provide a satisfying explanation for the target phenomena- they ‘bottom out’ (Woodward 2002). In contrast, when we conceive of a process, it is very often a process ‘all the way down’- perhaps to its most fundamental interactions, in many cases this will mean there is no natural place for an analysis of its dynamics to ‘turn its spade’ and stop. Consider the process of coastal erosion for example, a paradigm case of a process operating simultaneously at multiple levels of analysis. Large-scale phenomena such as currents and tides manipulate the oceans such that they interact more strongly with other large scale phenomena such as land-masses, at a level ‘down’, we may recognise waves as interacting

²⁶ In light of the phrase ‘screening-off’s’ established use (Brandon 1990, Sober 1992) we might be tempted - as we were with Dennett’s ‘real’ patterns - to use fresh and untainted terminology. Our use here however is quite in line with established use in some respects. We remain neutral here on whether the ‘screening-off’ of lower level events addresses the ‘levels of selection problem’ as suggested by Brandon, or if the suggestion that phenotypes ‘screen-off’ genotypes casually demonstrates natural selection acting on - and only on - phenotypes. We are partially committed to another of Brandon’s claims however: our use of ‘screening-off’ here may be ambivalent on *causality* but not on *explanatory clarity* and simplicity. When a mechanism screens-off lower level events, as with the omission of the movements of individual grains of sand in an explanation of coastal erosion, it will provide a better and more efficient explanation. Too much detail in this regard can compromise understanding, ‘screening-off’ provides a valuable service by ameliorating this.

with specific types of features like beaches, shorelines and cliffs- perhaps beaches with a flatter gradient are less prone to erosion while cliffs directly at the water's edge are more prone etc. At another level down again, we may be interested in the chemical composition of these features: is the process of erosion exacerbated by water with higher acidity? Does the composition, say the amount of limestone, of coastal features make their erosion more probable or more severe? An investigation of the process of coastal erosion undertaken at any of these levels would be warranted and potentially informative. This is not the case with mechanisms. They may share with processes a similar type of compositional hierarchy, but unlike those of processes some levels of mechanisms have clear *priority* over others²⁷.

In contrast, now consider a more traditionally-conceived mechanism: the machinery of a car plant. The mechanism of the car plant produces cars- this is the end product of the various sub-systems, operations or 'sub-mechanisms' that together comprise the plant's 'top' level mechanism. Notice that although the plant will no doubt be composed of various smaller sub-mechanisms, these do not contribute to the final product by producing *smaller cars* which then add up quantitatively to make normal-sized cars once production has ended. These sub-mechanisms produce car *parts* which are then jointly assembled to make the end product. Cars are produced at this final 'upper' level and only at this final level, they are not produced at any level below that. This dynamic is repeated in abstract form when parliament produces a new piece of legislation; the debates, the deals, the votes etc are all vital parts of the legislative mechanism, but they are not all acts of 'sub-legislation' which 'add up' to the final product. This was not the case for coastal erosion. At each compositional level in the processes described was erosion occurring: at the largest scale of oceans and land-masses, and at the level of beaches and waves, and at the level of the chemical composition of the water and coastal features. Different investigators of course may be interested in the different levels at which the process of erosion is occurring and may prioritise them accordingly. But this is up to them, their target phenomenon does not *naturally* provide any given level at which its investigation should be carried out, whereas an explanation in terms of mechanism

²⁷ As we saw in chapter 2's discussion concerning the viability of individual or population level causality, at which *precise* level the parts of the mechanism are identified may not always be obvious. For Walsh it was the interactions of individuals, for Millstein it was interactions of populations.

does have a natural, prioritised level: that one immediately below the level at which the target phenomenon is produced.

Should we wish to investigate the mechanism of our thermostat for example, we start with the target phenomenon of ‘temperature regulation’ and see how the component systems at the next level down together produce that phenomenon. This would no doubt include a thermometer and means of communication so that it can instruct the central heating as to what it should do, and both of these subsystems may be of interest in their own right- there is nothing stopping us from investigating further- but once we start investigating *them* our target phenomenon has changed, we will be looking at the mechanism of temperature detection and the mechanism of communication respectively. This ‘modular’ feature of mechanisms follows naturally from the observation that they are always mechanisms *of* some phenomenon; once mechanisms have their own sub-mechanisms they may well be ‘distant’ enough from the target phenomenon so that they are no longer mechanisms for that phenomenon. A sub-system that produces tyres is suitably distant from the target phenomenon of ‘car’ that they are not a mechanism of car production, and- lest we think this only applies to artifactual mechanisms- air humidity is suitably distant from the target phenomenon of ‘electric storm’ that it is not a mechanism of storms. Mechanisms bring to our attention the mutli-variate causes of phenomena but conveniently for our purposes also direct us towards the level of analysis that ultimately matters; the final level below which the phenomena of interest is produced.

Notice that mechanisms are also *regulatory* through their ability to reliably group entities and processes of the appropriate type. The elements that together comprise a pattern must of course share a similarity of some kind (only someone in the most severe grip of apophenia would claim to have discerned ‘patterns’ comprised from elements with no similarity whatsoever). But in many cases the precise nature of the similarity may not be readily apparent. A regularity of some sort might impress itself on us without us knowing explicitly what the relation between its component elements are or how they are achieved. As we will shortly see the converse is also true: some patterns may strike us quite forcefully as certain, objective and comprised from elements related in ways that are clear when in actuality they

are not. Suppose we have phenomena A and phenomena B, and we suspect that they are related in some way and that they may form part of a bigger pattern that we might be interested in investigating, but that we are unsure of- we just have a ‘hunch’ that A and B might be important in some way. We may look into the processes that cause A, and those that cause B and these may be suitably familiar enough for us to group the two phenomena together as elements of a pattern. The mechanism that produces the target phenomenon resists this kind analysis; we need a mechanism *for the pattern itself* not just for its elements. Consider again the example of Orion’s Belt. We may investigate the three stars that comprise the belt and find the mechanisms that firstly produced the stars and those that currently maintain them are very similar- near identical perhaps. Justifying our belief that these phenomena are of same type or kind. But these are production and maintenance mechanisms of stars, not of star-*constellations*. We require a mechanism to account for the target phenomena, the appearance of a belt-like formation as observed by someone on Earth, the mechanisms responsible for the existence of the stars themselves can be ‘screened off’ in this context. Although its elements and their relations are objective Orion’s Belt was not created through the actions of a mechanism, for this we would require a *constellation-creating* mechanism, which we do not have. Like the accidentally-produced Fibonacci sequence, we can provide no better explanation for its regularity other than saying its elements just happened to fall into that particular arrangement.

Even compelling patterns may be dismissed as impressionistic by virtue of a lack of mechanistic explanation. When one is available however, mechanisms may also be confirmatory, by showing that what may appear to be disparate phenomena are actually elements of a larger, though hard to discern, pattern. McShea & Simpson (2011, p22) provide an excellent example of this, which is made all the more compelling in that the phenomena it describes are not physical in the manner of constellations and coastal erosion but are historical incidents, illustrating that the principles of mechanistic explanation can be applied quite generally. They ask us to consider the following events, or in our parlance ‘elements’: the overthrow of the democratically elected government of Iran in 1953, the construction of the Berlin Wall in Germany in 1961, the invasion of Czechoslovakia by Soviet forces in 1968, the USA’s involvement in Vietnam from the 1960s to its withdrawal

in 1973. To this we may add the Moon landings in 1969 and Bobby Fischer's victory over Boris Spassky at the world chess championships in Reykjavík in 1972. This appears very much like the makings of an outlandish conspiracy theory; but these are in fact all elements of an authentic pattern that we can understand quite straightforwardly. They were all produced through the underlying common 'mechanism' of the Cold War that occurred between the Soviet Union and the USA in the later half of the twentieth century. Our impressions that these elements might be related in some way that is difficult to make fully explicit would be justified. The introduction of the mechanism confirms that these are elements of a larger pattern and that their relation is quite genuine independently of our perception.

Let us briefly recap the ground covered so far. We have discussed: *elements*- those parts of a pattern that together create the pattern's apparent features, these could be numbers in the case of the Fibonacci sequence or coastlines in the case of continental drift. We also identified: *relations*- the affinities between the elements of a pattern, these may be entirely impressionistic: like the 'belt-like' arrangement of Orion's belt, or genuine: like the ubiquity of DNA in terrestrial organisms. Lastly, we investigated *mechanisms* of patterns, noting that although the elements of a pattern may be the result of their own 'element mechanisms', patterns themselves must have their own 'pattern mechanism' if they are to be regarded as genuine. Some impressionistic patterns may be found to be less than fully real in virtue of this absence (Orion's Belt), while some initially quite dubious patterns (like the events underwritten by the Cold War) will enjoy strengthened credibility or even confirmation if the discovery of the right kind of underlying mechanism is found to be acting at the appropriate level.

With these terms in place let us turn our attention back to the different types of pattern that were identified by means of the three-way split established earlier in 4.2. Those that were originally filed under 'type ii)': patterns that were 'partially real' by virtue of their elements and relations being objective but whose generation we will now recognise as non-mechanistic we will term *quasi-patterns*. While those filed under 'type iii)': patterns which

featured objective elements and objective relations but whose generation was non-accidental and mechanistic we will term *authentic-patterns*.

At first pass, the distinction between quasi-patterns and authentic-patterns may seem fairly direct; quasi-patterns seem to be artificial and observer-dependent, whereas authentic-patterns are natural and independent of human observation. As we have seen, some cases slot neatly into this description: the Belt of Orion is dependent on observers like ourselves viewing its components from just the right perspective, while the pattern made by the ubiquity of DNA in living organisms exists whether we are observing it or not. But this classification is not quite so straightforward. Many, arguably most, of the patterns we experience in everyday life are quite authentic, but also ‘observer-dependent’ in the sense that we as observers were a necessary part of the process for their creation. Notice also that the patterns designed by human activity such as architecture, agriculture and other cultural manipulations of the natural world are dependent on their observers in quite a trivial way. It took human beings to *produce* these patterns; to create and arrange the pattern’s elements in their distinctive ways, but not to *maintain* these patterns; in that these patterns would still exist if human life were to disappear tomorrow afternoon. Now consider a subtlety-different case of observer-dependence: the patterns and regularities of economic activity. These are ‘observer-dependant’ like the artefacts of culture, in as much as they depend on the activities of human agents who also *happen* to be observers, but they are also objective in that they describe authentic patterns in collective human behaviour independently from any observation of them. The *role* of the observer is important here. Authentic-patterns may be natural or artificial, observed or unwitnessed, but they must not require observation to be authentic - consider again McShea & Simpson’s Cold War example.

The first moral to be gleaned from this extended section is that examples of both authentic- and quasi-patterns may present themselves to a number of different observers as being so clearly genuine that distinguishing between them may often prove impossible based on impressionistic information alone. This ambiguity works both ways. Not only can quasi- be mistaken for authentic-, conversely authentic- may be mistaken for quasi-; the fact that we are aware that our capacities for pattern recognition may slant towards the apophenic does

not mean that errors may not slip through despite this, even strikingly authentic patterns may be erroneously blamed on over-sensitivity²⁸. The second moral is that the discovery of a mechanism at the right level aids us greatly in deciding which type of pattern we are dealing with: by virtue of indicating which patterns are authentic- when a mechanism for it is found, and which may be quasi- when a mechanism is suspected to be absent. The third moral is that the complexity of a mechanism may often not be accurately reflected by the complexity of its description. Relatively simple mechanisms like carbon monoxide detectors, and very complex mechanisms such as those which produce storms can be described as though these phenomena were of comparable complexity when this is very clearly not the case in actuality.

With these morals fresh in our minds, let us see how they inform two particularly compelling natural examples. Being astronomical and geographical respectively, these continue our temporary hiatus from specifically biological and evolutionary phenomena a little further; it will be quite clear however that this is both relevant and informative.

4.3 Authentic & quasi patterns: two real world examples

From the perspective of an observer standing on Earth, the Moon and the Sun appear to be exactly the same size. This is demonstrated quite strikingly during a solar eclipse, the Moon obscures the Sun so precisely that aspects of the Sun's corona - usually swamped by the glare from the Sun's central body- can be witnessed creating the distinctive 'diamond ring' effect. If the Moon were too small the Sun would appear to us as a glowing donut, if too large both the Sun and its corona would be concealed. How do we explain this uncanny fit? We have *half* an answer. Their shared sphericity is due to the effects of gravity, the

²⁸ The twin dangers of mis-recognition can be straightforwardly understood through the following situation: We are lying in bed, semi-asleep, and hear unusual creaking sounds downstairs. We are alerted to the possibility that the noises may not be generated randomly by the house 'settling down' or similar, but are perhaps caused by an intruder and so constitute an alarming authentic-pattern... or perhaps the noises, although indistinguishable from those of an intruder, are entirely unrelated 'noises of the night' and so constitute a quasi-pattern? As indistinguishability works both ways, we might just as easily get up to confront an intruder when our alarm was actually caused by something benign and routine, as we might ignore the disturbance made by an intruder believing their movements to be random noise- especially should we believe ourselves to be somewhat apophenic and over-reactive.

material of both bodies are evenly attracted towards their centre of mass making their surfaces equidistant from this central point. There is of course an immense amount more to be said about how gravity produced and maintains the shapes of the Sun and the Moon, but the basic idea can be conveyed simply enough by the mechanism of a force which attracts matter equally through three dimensions. The specifics, the minutiae of their actual formation from gas and from rock can be ‘screened-off’. If it is just an explanation for their sphericity that we are concerned with, a simple mechanistic explanation concerning gravity and its general effects through the grouping of matter will likely prove sufficient. Including the extraneous detail may in fact be counter-productive by making the explanation *less* clear

What about the other half of the explanation? Like with the positioning of the stars in the Belt of Orion, we may be justified in thinking that the processes responsible for producing the Sun in its current size and relative position will be fully known and understood, and that the processes responsible for producing the Moon in its current size and relative position may be fully known and understood also²⁹. We are also justified in thinking that an explanation for why both bodies should appear to be the same size from the perspective of an Earth-based observer, *given* their actual sizes and positions, is available; the moon is around 400 times smaller than the sun and is also around 400 times closer. So it looks like we may have a complete, or at least satisfying, explanation for their striking similarity. We can explain the shapes of these two objects (gravity), we can explain why they appear to us as being the same size (distance in relation to actual size), and we are also confident that both the size and the distances of these bodies can be explained, in principle, by known scientific principles - or at least we have no grounds to suspect that this won’t be possible.

But clearly, this does not provide us with the right type of answer. We will either be discussing the specifics of optics and perspective for observers with our sensory capacities, or addressing two separate and distinct processes to account for each body’s individual formation- which resulted in their final products just happening to match up. What makes

²⁹ Although our moon is unusually large in comparison to the size of the planet that it orbits. There is ongoing debate as to precisely why this should be the case (<https://www.nature.com/collections/grlzpstdvb>)

either of these approaches unsatisfactory is that neither provide a mechanism for why both bodies should have formed in such a way and be placed in such a position, *such that* their sizes should appear to us to be the same. This seems to be a fairly clear example of a quasi-pattern, its elements are real, their relation is real (they do after all appear to us as the same size), but we do not have an underlying mechanism to explain the *pattern*, rather than its elements and positions. The closest that we could get to a ‘mechanistic’ explanation would strike us as being anything but- a vast collection of matter of various types in various positions and quantities and travelling at various velocities etc. There is a clear imbalance between phenomena and explanation here. The phenomenon could hardly be simpler: the precise match in size and shape between two simple and incredibly familiar objects, while anything approaching a satisfactory explanation for why they appear to us in this manner would be complex way past the point of plausible calculability. As there is no mechanism here to ‘screen off’ the lower level events responsible for their respective sizes we are forced to deal with them head-on.

Now to geology. Consider the case of continental drift; it had been noticed as early as the 16th century that the contours of South America’s eastern coastline matched closely to those of the western Africa, and by the 19th century geologists and naturalists were increasingly convinced that the similarities they were discovering in the fossil and sedimentary composition of the separated coastlines could not be coincidental (Oreskes 2018). In contravention to the prevailing geological orthodoxy at the time, the evidence appeared to indicate that Africa and South America comprised, either in whole or in part, a larger landmass that subsequently split. Perhaps surprisingly, it was not until the 1960s that this conjecture became scientifically respectable. The lack of any convincing mechanism which could explain the ancient separation of two continents by many thousands of kilometres meant that the apparent similarities in sedimentary composition and fossil type, and near-isometry in coastal shape were thought, on balance, to be more easily-explained away. This authentic- pattern was finally vindicated by the discovery of plate tectonics. By virtue of the pattern’s elements being the upper surfaces of enormous geological plates adrift on a sea of subterranean magma they are in constant, if unnoticeable, flux. This explained the matching patterns in coastal composition and shape. Initial impressions turned out to be correct; they

had not always been in their current locations and subsequent investigation suggested that previous to their separation approximately 145 million years ago Africa and South America had formed part of a single 'supercontinent' for hundreds of millions of years. The coastal edges of their surfaces were layered by the same combination and type of materials because it had accumulated in unison by the same sedimentary processes, and were inhabited by the same ancient species also, thus accounting for the similarities in the fossil record. What seemed very much like it could be a quasi-pattern was recognised to be an authentic-pattern through the identification of the pattern's underlying mechanism- quite literally underlying in this case. Although it now strikes us as undeniable, the match between the two coastlines by itself was not sufficient to justify the notion that their similarity was anything more than impressionistic, or at least coincidental, even when augmented by decades of additional circumstantial evidence. It was only by discovery of the mechanism responsible, quite invisible and distinct from the observed pattern, that justification for the pattern was made possible.

Both of these cases have much in common. They are probably the most widely-known patterns that are specific instances as opposed to everyday regularities, and the elements in both of these cases are very simple as indeed are their respective relations. Where they differ of course is in how we describe the causes of these patterns, one has mechanism to explain the similarity of their elements while the other does not. The Sun and the Moon's uncanny match remains 'one of those things' whereas South America and Africa's match is, if not comprehensively explained, at least explained to the extent that the pattern they comprise is considered authentic, rather than quasi by virtue of the mechanism responsible for it. That complex series of events may require complex explanations, and that simple series of events may require only simple ones will not strike us as particularly revelatory. The lateral division of a single object into two separate ones brought about by opposing forces is very straightforward and easily-conceived, as is an explanation for the similarity in their formerly-unified edges. We may demonstrate this to ourselves by ripping apart a piece of paper. Often the simplicity of the description accurately reflects the simplicity of the process, we have a simple phenomenon and a simple explanation to match. The explanation for the matching coastlines is very much of this type, as evidenced by its simplicity and

minimalism; we are able to briefly, but sufficiently, describe the processes responsible for their physical similarity quite straightforwardly. We could of course give a considerably more complex and detailed explanation for the geological processes responsible, but our knowledge of the mechanism means that this is an option, not a requirement that we are obliged to provide. This was also the case for our explanation for the shared *sphericity* between the Moon and the Sun; the economy in description matches appropriately the apparent simplicity of the process, we need only mention a very limited number of objects and their interaction according to easily relatable physical principles. The comparative *sizes* of the Moon and Sun on the other hand may appear to be a simple phenomenon but in actuality is not. It is the result of a very complex and disparate series of processes and events with no unifying mechanism. Without this there is no possibility of a mechanistic explanation, and without that there is no natural place for an analysis of the process to ‘bottom out’. We may imagine that a description, if one were possible, would have to be highly complex and convoluted and so reflect reality of the process itself by being so.

4.4 Wilson and Dawkins: authentic or quasi?

With the above contrast between the authentic-pattern produced by the Atlantic coastlines the quasi-pattern produced by the Sun and Moon in sharp focus, let us return to evolutionary biology and recall the patterns offered by Wilson and Dawkins at the start of this chapter: various impressionistically ‘advanced’ phenotypic and behavioural traits and abilities showcasing evolutionary advance, and adaptive trajectories respectively. Are these patterns authentic like the coastlines of Africa and South America, or are they quasi like the sizes of the Sun and the Moon?

Following our analysis above we now have a five-part heuristic to help us find out: i) the patterns should not be accidental, but produced through an underlying mechanism; ii) this mechanism should be reliable and repeatable- if not repeated; iii) the mechanism should simplify the pattern’s explanation by ‘screening off’ lower-level events when required, so providing a ‘natural level’ at which their explanation should be provided; iv) the mechanism

should be regulatory, in that it should indicate which elements belong to the pattern in question, and which ones do not; and finally v) the mechanism must be a mechanism for the pattern in question, its elements and their relation- explaining the production of the elements alone without accounting for their integration is insufficient.

Although our focus will be on Dawkins' and Wilson's evolutionary patterns out of the many available they are in fact representative of two quite different ways of conceiving of progress in evolution. At one extreme we have the traits and abilities of individual organisms; how they compare with each other, how they enhance fitness, how well they perform their biological functions; we can see our limited defence of normative language to describe performances in biological function (in Chapter 3) falls squarely within this 'parochial' perspective. At the other extreme we have evolutionary phenomena at the largest and grandest of scales; is there some way that life taken as a whole has 'progressed' since its inception? Are organisms taken from later epochs qualitatively different in some way to those that preceded them? Clearly there is no theoretical restriction which precludes holding both of these views simultaneously and indeed Dawkins in places appears to do so (Dawkins 1996a & 2003, Ruse 2006), although we will concentrate on his ideas concerning smaller-scale patterns for the remainder of this chapter.

We will start with Dawkins, for whom natural selection's 'progressive' capacities are a *theoretical* necessity- a general claim quite separable from any particular evolutionary pattern that its actions might produce. As long as we accept that there are traits which require multiple generational stages for their gradual construction, and as long as we hold the retained selective benefits of random variation at each stage to be responsible, we must also accept that any 'multi-stage' trait - by virtue of its existence alone - acts as evidence for evolution's 'progressive' capacities. Following on from our observations in chapters 2 and 3, Dawkins' perspective here may strike us as somewhat familiar. In chapter 2 our defence of natural selection's creative abilities was based upon its ability to vastly increase the probability of traits requiring multiple and collectively-unlikely combinations coming about, whereas Dawkins' emphasis is on the requirement of such a mechanism for the incremental evolution of complex traits. The traits which we considered to be examples of extended

progressive evolution in chapter 3 and the traits which Dawkins considers to examples of ‘progressive’ evolution would seem - impressionistically at least - to coincide also:

“the evolution of echolocation in bats and river dolphins.... the evolution of electrolocation in fish... skull dislocation in snakes for swallowing large prey... the evolution of the complex adaptations that equips cheetahs to kill, and the corresponding complex that equips gazelles to escape” (Dawkins 1996a, p1018)

One in particular- the evolution of the mammalian eye- which Dawkins addresses in both (1996a) and at length in a separate work (Dawkins 1997) will also be one of our central concerns in Chapter 5. If Dawkins has commented explicitly on the SE conception of biological function he has done so discretely, although we might imagine someone with his established record of adaptationism would very likely give it sympathetic hearing. Notice the emphasis on the *functions* of the adaptive structures in the passage above:

‘echolocation’; ‘electrolocation’; ‘swallowing’; ‘equips... to kill’; ‘equips... to escape’. It is not cumulative selection *per se* which Dawkins considers progressive, but in common with our findings in chapter 3, it would seem to be cumulative adaptation producing functional improvements. There is evidently much overlap between our views and Dawkins’ in this area.

But notice that in addition to Dawkins’ theoretical claim - which we might agree with - there is also an empirical one which is more controversial. Dawkins commits his brand of progressive evolution being present in the *majority* of evolutionary lineages; essentially a restatement of ‘empirical adaptationism’ (Godfrey-Smith 2001) which considers natural selection to have been by far the most significant influence for the majority of evolutionary phenomena - and for which Dawkins’ enthusiasm is renowned. The long running debate concerning adaptationism in all of its varieties remains contentious (Gould & Lewontin 1979, Dawkins 1983, Orzack & Sober 1994, Dennett 1995b, Godfrey-Smith 1999 & 2001, Sandson 2003, Lewens 2009), and we offer no stance on this here. We did suggest however in chapter 3 that a considerable amount of natural selection will be due to improvements in functional performance, rather than idiosyncratic or novel effects - a claim concerning how

much of natural selection is progressive in comparison to non-progressive natural selection; this is in clear contrast to Dawkins' claim which concerns how many evolutionary lineages feature progressive natural selection in comparison to lineages which do not. We are committed to most natural selection being progressive, Dawkins would seem to be committed to most *evolution* being progressive - a much stronger and wider claim, and one that we may be somewhat reluctant to support. Let us however interpret Dawkins' stance here more generously, and suggest that many - if not all - evolutionary lineages show the type of 'adaptive progress' that he describes. How do the patterns of these organisms' gradual adaptation to the demands of their environments fare in the light of our five-part heuristic?

We can be confident that this progressively-adaptive 'fit' of organisms is *not accidental*; we have a clear and well-understood mechanism - natural selection - to account for this phenomenon. The echolocation of bats and dolphins for example are traits which have expressly evolved via the mechanism of natural selection in response to the nature of the environments in which the ancestors of these creatures inhabited. But is the mechanism responsible for these phenomena *reliable* in the sense that given very similar initial conditions, the same selection pressures, the same probabilities for the arrival of the required variation etc, we could expect it to achieve similar results? Although the debate surrounding the contingency of evolutionary phenomena continues (Gould 1990, Beatty 1995 & 2006, Losos et al. 1998, Conway-Morris 2003, 2006, 2008 (eds.) & 2010, Vermeij 2006, Arendt & Reznick 2008, Powell 2012, Blount 2016, Blount, Lenski & Losos 2018) it is clear that natural selection is not merely potentially repeatable, but has actually been repeated many times over evolutionary history. Even if we consider these convergences to be of less evolutionary importance than some authors suggest, it is clear that the mechanism of natural selection does not operate on a random, chaotic basis. Selection may be probabilistic, but it is not arbitrary, and when we see natural selection promoting one characteristic at the expense of another, empirical investigation can often find out why (Reznick 1985, Endler 1986, Barros 2008, Sober 2008). If natural selection did not operate in a structured and reliable manner investigations like this would not be possible.

The mechanism responsible for Dawkins' patterns provides us with a mechanistic explanation, and this conforms neatly with the third heuristic by '*screening-off*' of the lower level events underlying the surface phenomena also. If we wanted to understand the progressive adaptation of certain snakes' capacity to unhinge their jaws for example, we do not need to know the precise ancestral population dynamics which gave the earlier versions of this trait a selective advantage, nor do we need to know the exact genetic mutations underpinning the traits' variation or how they interacted with the organisms' developmental systems. Our knowledge of the mechanism allows us to ignore these details and concentrate instead on the phenomena that they are responsible for producing.

The mechanism of natural selection also acts to *regulate* the suitability of elements in the patterns that Dawkins describes. Suppose we wished to verify ancestor-descendant samples in an aquatic evolutionary lineage charting its members' progressive development of electrolocation. Knowing that the mechanism responsible for the trait operates through the retention of selectively-beneficial and gradual variation, we would not only very likely be able to discern with a high level of accuracy if any given sample belonged the sequence, we would also likely be able to arrange the samples which did belong in their correct temporal and evolutionary order, so the desideratum that an authentic pattern should have a mechanism which *accounts for pattern itself* - both its elements and their relations- would also seem to be sufficiently satisfied. In fact the relationship between the cheetah and the gazelle arm's race in Dawkins' final example reflects the isometry of the South American and African coastlines in some respects. Clearly there are differences, we would have to stretch very far back into evolutionary history to find a common-ancestor which then 'split' to form the cheetah and gazelle lineages, and unlike the coastlines on either side of the Atlantic the 'fit' of each lineages' predator and prey adaptations was continually shaped by forces occurring long after their original 'splitting event' occurred. However, in both cases we account for the complementariness of the patterns' elements through their synchronised production by a common underlying mechanism, and in both cases each side heavily influences the shape of the other in accordance with the mechanism's own dynamics.

We might suspect that arms-races' enduring appeal for Dawkins (Dawkins & Krebs 1979, Dawkins 1986, 1996a, 1996b & 2004) may be again due in part to his adaptationism. The importance of inert and potentially capricious environmental variables in an account of adaptation by natural selection are much reduced if the animate and, importantly, co-evolving environmental variables of a population's prey or predator species are given preferential focus; an observation that might be only of marginal interest if arms-races did not also provide Dawkins with a means of characterising adaptation as *improvement*. This again will strike us as familiar: in chapter 3 we saw the Selected Effects approach was able to provide an objective basis for comparing earlier and later traits' functional performances, here arms- races perform a distinct although similar justificatory task. In the paradigm arms-races that Dawkins refers to (Dawkins 1986) later prey are able to avoid or defend against earlier predators more effectively than their ancestors, as being able to do this was the primary means through which their relative fitness was raised. Equally, later predators become more effective at catching and killing their prey as this was the primary means through which their relative fitness was improved. Here we see improvement not by comparing directly the performances of descendant organisms with those of their *ancestors*- as in chapter 3, but by comparing descendent organism's performances with the performances of their *ancestors' predators or prey*. The directional adaptation of each warring lineage as their arms-race escalates should - *ceteris paribus* - promote this. As far as one species comprises an environmental variable for another, Dawkins reasoning here is sound, although how prevalent and important an influence the arms-race dynamic has been in evolutionary history is debatable (Parker 1983, Thomson 1986, Brodie III & Brodie Jr. 1999). We might certainly venture that they would not be characteristic of the 'majority of evolutionary lineages' and can be confident that they will be more exceptional than the functional improvements which we surveyed in chapter 3. Indeed, arms-races - where a species' functional improvements are matched by another species' which in part comprise the former's environment- are a special case of this.

Dawkins' particular brand of adaptationist progressiveness then is grounded in patterns which, according to our heuristic, appear authentic. Their generation is not accidental, the mechanism which produces them is sufficiently reliable, a description of the mechanism's

operation ‘screens-off’ sub-mechanism processes and events, the mechanism can regulate which phenomena belong to any particular pattern and which do not, and the mechanism is able to provide a comprehensive explanation for the patterns it produces- both their elements and their respective relations. Additionally, Dawkins also provides a means through which cumulative adaptation may make ‘progress’ as well as being ‘progressive’. The particular evolutionary dynamics of arms-races, albeit heavily idealised, suggest that descendant organisms will be better-equipped to catch their ancestral prey or avoid their ancestral predators.

We now turn to progress’ second defendant: E.O.Wilson. Despite recent antipathy (Dawkins 2012, D.S.Wilson 2012, Johnston 2014) Wilson and Dawkins share an enduring belief that progress in evolution is real, and both contend that it is not merely contingent or incidental but a core characteristic of how evolution on earth has unfolded. Wilson’s views differ from Dawkins’ on a number of fronts. Firstly, the patterns that Wilson suggests demonstrate evolutionary progress are firmly manifest at the largest and grandest of scales and are a “property of the evolution of life as a whole” rather than a phenomenon relatively localised to natural selection operating within lineages; secondly, Wilson does not provide us with any obvious mechanism to account for his expansive list of ‘progressive’ biological phenomena; although thirdly, does provide guidance as to the effects that progressive evolution has had on the features that he lists as examples: body size, brain and behavioural complexity, precision of environmental control: ‘... the overall averages of these traits and their extremes went up.....’ (Wilson 1991, p175).

Elsewhere (ibid p174) Wilson emphasises that his perspective on evolutionary progress is not just expansive in the taxonomical sense but expansive temporally, extending back into the distant past to the beginning of life itself. Perhaps surprisingly for a supporter of evolutionary gradualism (ibid p80)³⁰. Wilson also contends that this history of evolutionary progress can be marked out by ‘four great steps’:

³⁰ Indeed, Dawkins himself suggests that certain innovations including the eukaryotic cell may have been ‘watershed events’ in the history of life (Dawkins 1996a, p1019) despite being an avowed gradualist.

1. Life's inception 3.9-3.8 billion years ago
2. The origin of eukaryotic organisms
3. The Cambrian explosion
4. The origin of the human mind.

Readers familiar with the suggestion that the history of evolution can be partitioned in this way may be reminded of Maynard-Smith and Szathmary's 'Major Transitions in Evolution' (Maynard-Smith & Szathmary 1995). Both patterns have much in common: they start from life's inception, terminate with the advent of human social intelligence, place the advent of eukaryotic organisms somewhere in the middle, and organise the history of evolution into distinct stages characterised according to the latest 'evolutionary innovation'.³¹ Additionally, the events that Wilson contends mark-out the history of progressive evolution also share the Major Transition's increasingly-recognised difficulties with consistency (Calcott & Sterelny (eds.) 2011, Szathmary 2015, O'Malley & Powell 2016), and to this we may add that both Maynard-Smith & Szathmary's and Wilson's patterns share difficulties with their authenticity.

Consider Wilson's series which charts the earliest and simplest forms of life to the origin of the human mind and the intermediate era-defining events. Clearly for life to evolve, life first has to start, so the origin of life is not strictly an event *in* evolutionary history whether we think the patterns that it subsequently gave rise to are authentic, progressive or otherwise. This observation may strike us as somewhat pedantic, but notice that by starting from this earliest of 'life history' events Wilson thereby includes an *accidental* occurrence as one of his larger pattern's elements, something which Dawkins through his focus on the traits of already-existent organisms avoids. Shouldn't patterns of evolutionary events at minimum commence with those brought about by *evolutionary* processes? There is no easy answer to this. For natural selection to happen there has to already be a population of variants whose reproductive success is influenced by their variation, and whose offspring are more similar

³¹ In fairness, Maynard-Smith in the paper-length precursor to the Major Transitions (Maynard-Smith 1988) explicitly distances himself from idea that the stages described amount to *progress* of any kind, a statement which is later reiterated in his co-authored book-length account. There is little no doubt however that both versions of the Transitions consider the evolutionary patterns they describe as *authentic*.

to themselves than to the population average - a significantly complex arrangement in itself, and one whose origins we are justifiably curious about. Many 'histories of evolution' like Wilson's start with the beginnings of *life* rather than the beginnings of *evolution* so if this approach is a mistake it is not Wilson's alone, but notice that if Wilson had started from the earliest evolution this would not only make his first stage more thematically consistent with his later events, it would also greatly reduce its accidental nature. As we recognised in chapter 2 and reiterated earlier in this chapter, evolution by natural selection may not be determinate but it is not accidental.

Even granting that this correction is one that Wilson would tolerate - or is one that we are paying undue attention to - the second stage: 'the origin of eukaryotic organisms' suffers from a similar difficulty. Despite its undoubted significance, the belief that the endosymbiotic event which created the first eukaryote represents a paradigm evolutionary accident is almost unanimously-held in evolutionary biology (Gould 1990, Gray et al. 1999, Embley & Martin 2006, Koonin 2010, O'Malley 2010, Archibald 2015). If orthodoxy is correct and this was a 'freak' event, it is difficult to see how it could be incorporated into anything resembling a wider authentic pattern which would require other elements which resemble it. Now perhaps this is unfair for couple of reasons. As mutation and - to a lesser extent - recombination are the ultimate source of biological novelty, surely *any* evolutionary innovation will be accidental when it first occurs. Isn't it what these events lead to, rather than their origins what is important here, and what Wilson is drawing our attention to? This defence is, I think, only partly successful. A relevant exception is the novelty created when unrelated lineages combine to the extent that we consider their joint product to be a single entity. This is not the case in parasitism or almost all symbiosis - where the meeting of separate lineages involved may result in highly co-dependent but sufficiently distinct organisms - but has occurred in the history of evolution at least twice: in the events which created chloroplast organelles and the events which created eukaryotic cells. In neither of these cases was the resulting novelty caused by mutation or recombination, but by the coming together of very different organisms to make an entity quite distinct from either. So while it is true that eukaryotes (and chloroplasts) in common with all (non-synthetic) biological novelty are the results of accidents *per se*, their specific, idiosyncratic nature

might perhaps give us pause that they should not be incorporated as elements in a larger pattern without first taking this into account.

But suppose we are unconvinced by this line of argument and believe that the causes of certain events don't matter so much as their later consequences? Perhaps the accidental non-biological nature of the origin of life and the accidental and somewhat esoteric nature of the first eukaryotic cell can be overlooked in the light of the evolutionary developments that these events made possible? This may appear to be a viable approach concerning the origin of life which made *all* subsequent evolutionary events possible, and to a lesser extent for the advent of the eukaryotic cell which appears to be a necessary precursor to the sort of evolutionary phenomena which Wilson considers progressive: '...feeding and defensive techniques, brain and behavioural complexity, social organisation' (Wilson 2001, p175). This ability to provide a platform for further evolutionary development is also characteristic of Wilson's third stage- the Cambrian explosion, so we might be forgiven for thinking this might provide a unifying theme although in contrast to the origin of life and the advent of the eukaryotic cell there is significant doubt that the Cambrian explosion actually happened- at least as traditionally conceived (Runnegar 1982, Wray et al. 1996, Conway-Morris 2000, Fortey 2001, Butterfield 2003, Briggs & Fortey 2005, Priest 2017). If the understanding that something quite notable and enduring occurred to biological diversity at this time is approximately correct however, then including this stage alongside the origin of life and the arrival of the eukaryotic cell might seem to have some justification. But notice that this quality of promoting further evolutionary diversification is noticeably absent from the final stage- the origin of the human mind. It seems far from obvious that the arrival of human intelligence has provided a platform for further evolutionary developments (artificial biology notwithstanding). The arrival of *Homo sapiens* has diversified human beings, both genetically and geographically but, if anything, has reduced net evolutionary possibility by its contribution to the extinction of congeneric species (Banks et al. 2008, Callaway 2016, Varki 2016) and more recently to biodiversity more generally (Carpenter & Bishop 2009, Yule et al. 2013).

So far Wilson's pattern of evolutionary progress does not hold up to scrutiny. The causes of the pattern's elements are in part accidental, the reasons for including each of these stages as elements is unclear, and the esoteric nature of eukaryotic cell in particular certainly goes against the grain of our contention that the generation of a pattern should be *reliable*, and we seem to be very far from being able to provide the required *mechanism for the pattern*. Notice also the idiosyncratic nature of the origin of life and the advent of eukaryotic organisms not only makes them accidental, it additionally prevents us from *screening-off* the minutiae of their specific occurrences. Any satisfactory explanation for these elements would be forced to deal with the specifics of pre-biotic chemistry and cell biology respectively because it is at these levels of analysis that the events occurred, they only become evolutionary important by their subsequent effects and even this loose description does not apply to the final stage of human intelligence the effects of which are highly questionable.

But perhaps most damagingly, the pattern described by Wilson does not provide any *regulatory* role. Biodiversity is clearly important to Wilson, which has been: '...pulled along by evolutionary progress, measured in four great steps' (ibid p174), but it seems cannot be equated with progress directly as: 'Biological diversity embraces a vast number of conditions that range from the simple to the complex, with the simple appearing first in evolution and the more complex later' (ibid p175), while: '*an undeniable trend of progressive evolution has been the growth of biodiversity by increasing command of the earth's environment*' (ibid p175, italics added). So if progress is not made by expanding the arena of evolutionary *possibility*, maybe *actual* biodiversity has a lot to do with it. But as we are not given any consistent rationale why those four elements in particular were included we lack the reasons for the exclusion of additional elements which Wilson misses out and would- at least intuitively- seem to augment Wilson's perspective of natural history quite well. Consider the following seconded from Maynard-Smith & Szathmary's Major Transitions (1995 p6): 'Asexual clones to Sexual populations'. If evolutionary progress is a cause of ('pulling biodiversity along'), a consequence of ('a property of the evolution of life') or simply *is* biodiversity then this particular evolutionary stage has equal if not arguably more qualification for inclusion in Wilson's four great steps than those originally

chosen; consistency would seem to demand its inclusion. The advent of sexual reproduction not only brought with it a doubling of local diversity within those species which additionally developed 'specialised' female and male variants, the consequences of meiosis and genetic recombination near-guarantee that each individual organism within the species will be genetically unique, even from close relatives. More importantly, the reproductive barriers thrown up by genetic diversification in formerly-compatible sexual populations are a well-established and effective means of entrenching incipient speciation, and by extension promoting the global biodiversity apparently unpinning Wilson's entire project. If biodiversity has relevance to the evolution's 'great steps' sexual reproduction's exclusion here is genuinely puzzling.

So even a liberal interpretation which viewed the progressive steps of evolution as having *some relation* to increasing biodiversity and/or further evolutionary development would appear to require substantial revision. Wilson's first three stages fit either of these characterisations reasonably well, but the fourth - the advent of the human mind- has likely decreased biodiversity, and offers no clear suggestion as to how its arrival has opened up new evolutionary possibilities. Additionally, if either of these characterisations are approximately representative of what Wilson believes evolutionary progress consists in, then the pattern of progress as charted by his four steps is incomplete, there are good reasons to think that sexual reproduction should be included as a subsidiary step and every possibly that other events (cell differentiation? the establishment of eusocial colonies?) should be included too.

We started this chapter by observing the general characteristics of natural patterns, and drew a line between patterns which were authentic and patterns which were quasi. As demonstrated by the sizes of the Sun and the Moon, how compellingly a pattern may strike us is no guarantee of the pattern's authenticity; as demonstrated by the near-isometry between the Atlantic coastlines, the discovery of a mechanism for the pattern may prove definitive in favour of its authenticity. Authentic patterns should be: non-accidental, reliable, screen-off lower level events, regulate the appropriate included elements, and finally

account for the relations holding between the pattern's elements. Dawkins' pattern satisfies these criteria, whereas Wilson's pattern does not.

4.5 Higher-order mechanisms: some desiderata

It might be tempting to think that the reasons for Wilson's pattern lacking authenticity belonged to Wilson alone, and that the stages he believes are characteristic of evolution in the long term are either wrong, incomplete, or were never designed with the above philosophical scrutiny in mind. There is a strong case to be made for each of these, but in as much as Wilson's stages are generally representative of evolutionary patterns at the largest of taxonomic and temporal scales, his difficulties are also theirs. Consider what it would mean for a pattern in the history of life to be considered authentic. Firstly, the generation of each of the pattern's elements should be *non-accidental* by virtue of production through an underlying mechanism. This was not a problem for Dawkins as this desideratum was covered by natural selection, but at scales above intra- and perhaps inter- species selection the effects of natural selection are questionable if not outright denied (Lewontin 1970, Jablonski 1986, Gould 2002, Simpson 2010, Okasha 2012a, Hoehn et al. 2016). If natural selection cannot be relied upon as the mechanism to explain regularities at these higher taxonomic scales we require an alternative one that can, and these are in short supply.

It has long been thought for example that life's 'complexity' has increased over evolutionary time, both in the sense of increasing 'disparity' of differentiated anatomy in individual organisms and in the 'diversification' of species and higher taxa more broadly (Lamarck 1809, Spencer 1891, Gould 1998, Adami et al. 2000, McShea & Brandon 2010). What mechanism do we think could account for this? The answer is often presented through a combination of the following principles: the first life was by necessity simple, there is 'left-wall' of simplicity below which life processes will not function making 'increased complexity' the only option, and that any group of homogenous entities will tend to differ over time as minute differences in their interactions mount up. It is admittedly difficult to see how these principles might not have had some influence on the patterns of increasing

biological complexity- in so far as these patterns actually are realised empirically- and so to some extent provide the required mechanism. But there is an issue here which should temper our enthusiasm: the ‘mechanism’ is only *incidentally* biological. Assuming a lower limit on difference we could just as easily apply these principles of change to any entities upon which differences accumulate, these could be billiard balls, twins, or pebbles on a beach arranged according to Dennett’s real pattern (Dennett 1991) - indeed one of Gould’s well-known criticisms that this evolutionary trend could be revelatory in anyway likened the effects of a ‘fixed floor’ coupled with a change in variance like this to the scoring of professional baseball and the meandering of a drunkard’s walk (Gould 1998). The distinction is a subtle one: by virtue of the ‘complexifying mechanism’, there could be grounds to claim that there are authentic patterns *in* evolutionary history- the elements are biologically instantiated after all- but this falls short of the stronger claim that these are authentic patterns *of* evolutionary history, when the mechanism’s non-evolutionary effects can be witnessed elsewhere. For this stronger claim we would need a distinctly *evolutionary* mechanism to produce the higher taxonomic phenomena in a similar way to which natural selection provides an explanatory mechanism for the ‘lower level’ adaptationist patterns suggested by Dawkins. Notice that this is not just an issue for Wilson’s increasing biodiversity or for the separate claims of increasing complexity, but will be a desideratum for *any* evolutionary pattern discerned in taxa from genera and above. Aside from the highly speculative suggestion that this ‘higher order’ mechanism could be natural selection itself a candidate has yet to appear.

Suppose however that we are comfortable with an explicitly biological pattern’s authenticity relying on a mechanism that is only incidentally biological, and are persuaded that either the above ‘complexifying mechanism’ or something similar might be adequate to explain evolutionary patterns found at high taxonomic scales. The mechanism would have to be *reliable*. The ‘complexifying mechanism’ for example seems to comfortably provide this; given cumulative variation and a ‘floor effect’ McShea & Brandon point out that we can quite reliably expect heterogeneity to increase- and, if anything, may need to find reasons for diversity and disparity *not* being higher than is currently observed (McShea & Brandon 2010). The mechanism would also have to be *regulatory*, but given that the pattern is the

entire living world deciding which elements are fit for inclusion might be the easiest of the five desiderata to satisfy.

But notice the mechanism would also need to ‘screen-off’ lower level effects, and while we cannot tell whether an as-yet-unidentified mechanism has the potential to do this, the ‘complexifying mechanism’ runs into profound difficulties here. As this mechanism is not exclusively evolutionary or biological but physical we have no reason to stop an investigation into the mechanism’s operation at the evolutionary or biological patterns that the mechanism produces. Higher taxa will be diversifying due to the influence of this mechanism but not only higher taxa - so will species, organisms, and parts of organisms like cells and their organelles, as well as mutating genes down to the molecular level. This seems very unlike a mechanism to explain diversity at level of say genera or familiae, the way that natural selection explains the adaptation of populations or the improvement of biological functions and much more like the processes of coastal erosion, wearing down cliffs, stones, pebbles and sand. There is no preferential place in either of these compositional hierarchies at which these processes ‘act’, and so no principled means to decide at what level an analysis should be undertaken. Again, this may not bother us, but if it is ‘large scale’ evolutionary patterns that we are interested in, the ‘complexifying mechanism’ can tell us nothing more about them. The apparent rises in anatomical disparity and ecological diversity become just ‘one of those things’ much as a later coastline might be expected to be increasingly worn away given enough tide and time. At best the gradual creep of increasing heterogeneity might make for an important background condition which should be factored into the calculations of the effects of *other* mechanisms- as we might factor-in drift when calculating the strength of natural selection, or wind speed when flying a plane- but this is a long way from providing anything like a satisfactory explanation for the target phenomena.

Lastly the mechanism would have to be a mechanism for both the pattern’s elements and their relation. For patterns found at higher taxonomic scales, how effectively a mechanism is able to account for this depends of course on what phenomena are grouped together to make an ‘element’ of the pattern. As there is no natural place where ‘screening-off’ can be carried out this may be done in a number of ways. We might for example find a pattern in the living

world at the level of genera, or we might see a pattern composed by connections holding between classes or familiae, in which case the pattern's elements will be genera, classes and familiae respectively. But there is plenty of potential here for the elements of a pattern to be composed from phenomena which do not fit our established means of biological classification quite so neatly. Without the procedure of screening-off to provide the appropriate level, the identification of elements becomes something of a free for all. A pattern might present itself through the use of para-phyletic elements such as Reptilia or Pisces, or consist of elements comprised with little regard for established classification at all: perhaps 'counter-shaded aquatic predators' or 'plants with multi-tipped leaves'. The mechanism responsible for the pattern's generation should account for the relations which hold *between* such elements, and while again we cannot predict the effectiveness of a mechanism we do not know, the 'complexifying mechanism' fares poorly here. The 'instability of the homologous' may account for increasing differentiation *within* a given element, but will not account for any differences in the rate, kind or extent of differentiation *between* the selected elements that together comprise a higher order pattern.

Let us unpack this claim. The complexifying mechanism for example might explain the increasing disparity and diversity found in tropical regions, and we might take this grouping as the basis for an element in a wider pattern; the complexifying mechanism might also explain the increasing disparity and diversity found in the polar regions and use this as an element also. But what the complexifying mechanism cannot do- and what *any* mechanism that purports to explain large scale evolutionary patterns must do- is to explain how the elements of the large scale pattern relate to *each other*. In this specific example it would have to give reasons for why the disparity and diversity in one geographic region is different in kind, rate or extent from those in another. But as we saw clearly in our example of the Sun and the Moon's similarity, it is often inadequate for two elements be a product of a given mechanism. Gravitational forces may account for their spherical shapes, just as the complexifying mechanism accounts for increasing differences at both geographical regions, but if the pattern is to be authentic the mechanism must also account for why one element is bigger, smaller, comparable, faster or slower etc than the others which make up the larger pattern. Gravity alone does not do this for the similarity of the Sun and the Moon, and the

complexifying mechanism does not do this for the variable rates of disparity and diversity in our different geographic regions.

We should of course not preclude the possibility that a mechanism which can sufficiently account for large-scale evolutionary patterns could be found. The analysis undertaken in this section is not intended to be moratorium, but instead an indicator of the sort of properties that a putative evolutionary mechanism should possess in order to make the patterns it creates authentic. As we saw in Dawkins' examples - and recognised in our own - natural selection does make some of evolutionary patterns at the population and species level authentic, if its dynamics *can* be reproduced above these higher taxonomic levels, perhaps any patterns that it produces there could be similarly authenticated. But whether the mechanism is natural selection, is partly natural selection, or is a mechanism entirely different to natural selection, our work in this chapter has clarified the qualities that it should possess, it should be: non-accidental, reliable, screen-off lower level events, regulate the appropriate included elements, and finally account for the relations holding between the pattern's elements. Currently, there are no obvious candidates.

4.6 Chapter recap and conclusions

In section 4.2 we looked at the general properties of patterns and identified the different types that our observations of nature present to us: 'authentic-patterns' and 'quasi-patterns'. We observed the role that a mechanism plays in discriminating between these different types of pattern and recognised a five-part heuristic through which authentic or quasi patterns could be recognised. In 4.3 we addressed two real-world natural patterns to demonstrate the practical application of these heuristics: the close similarity of the Atlantic coastlines was authentic according to this heuristic, whereas the sizes of the Sun and the Moon were found to be quasi. In 4.4 we addressed the respective patterns suggested by Wilson and Dawkins to be present in the history of life. Here again the application of our heuristic proved discriminatory: we observed that Dawkins' suggested pattern was authentic, while Wilson's was not. Then in 4.5 we discussed the possible ramifications of finding what initially

appeared to be an authentic pattern in ‘large-scale’ evolution to actually be quasi-, and suggest the properties that a genuinely authentic pattern at taxonomic scales at genera and above would have to possess.

As quasi-patterns can be as, if not more, compelling than authentic-patterns our impressions that a pattern is genuinely ‘out there’ and not a product of our natural inclination to find order and regularity in disparate phenomena may often mislead. While events in the natural world have causes, not all patterns in the natural world have explanations, and to ensure an investigation is not an attempt to join up the random dots of unrelated elements, the discovery of a mechanism for the perceived pattern will be of significant help. This will be the case whether the elements of the patterns are electronically generated numbers, geological structures, planetary/stellar phenomena or phenomena recognised through evolutionary or biological science. The patterns of gradual adaptation proposed by Dawkins to be present in the ‘majority’ of evolutionary lineages are perhaps not as endemic to evolution as Dawkins suggests, although the progressive type of natural selection we identified in chapter 3 proved sufficient explanation for the examples provided. We did not rule-out the discovery of a mechanism which could explain large-scale evolutionary patterns such as the one suggested by Wilson, but equally - given the requirements brought to light our five part heuristic - we may be justified in being sceptical as to its discovery.

In chapter 5 we draw on the previous chapters’ various findings in order to analyse further: the progressive evolution brought about by natural selection’s causality and creativity - as argued in chapter 2; incorporate into our analysis the improvements in biological function - as identified in chapter 3; and respect the criteria for natural patterns - identified earlier in this chapter - that sustained progressive evolution can produce. As we noted in our introductory chapter, perspectives regarding progressive evolution are often accompanied by anthropocentrism to their detriment. In chapter 5 we face this issue head on, and find that the link between these two positions is weaker than supposed; we may have one without the other, and the undoubted tendency to think of human beings as superior from an evolutionary perspective is not sufficient grounds for us to think it untrue. When then look at the evolution which occurs within lineages. All species which currently exist are the most

recent members of lineages of descent reaching back to the beginnings of life itself, although thinking of evolution in this way has similarly been treated with suspicion. Not all of this is unwarranted, our representation of evolution within lineages over both longer and shorter timescales can often mislead us as to the reality of the evolution depicted. Thus far our commitment to the SE conception of function has limited our scope of functional comparisons to traits belonging to organisms of the same species, but due to the sometimes radical changes which species undergo over longer spans of evolutionary time our basis for functional comparison needs further amendment. Progressive evolution within lineages over the longer term requires special handling, we observe the nature of this in the chapter ahead.

Chapter 5

Progressive evolution within lineages: representation and reality

5.1 Introduction.

Although we saw in the previous chapter that the identification of an underlying mechanism provides sufficient warrant to take a pattern's authenticity seriously, we have additional *prima facie* reasons to consider the regularities perceived in evolutionary lineages to be authentic. Evolutionary lineages are objective and not merely imposed on the phenomena by the prejudices of observers. Even when separated by vast spans of time and substantial differences in function, trait, or genotype, the properties of ancestral lineage members have direct causal and explanatory relevance for the properties of their descendants; inheritance through descent guarantees this. Resemblances between earlier and later members of an evolutionary lineage will not be accidental, but due to inherited similarities in genotype and developmental mechanism. By virtue of the objective nature of the relationships which hold between their component organisms, evolutionary lineages are also regulatory: it will be matter of fact for any given lineage whether an organism belongs to that lineage or does not, however difficult this might be to establish in fact.

However as we also saw in chapter 4, having elements linked by an ancestor-descendant relationship is not by itself sufficient to establish a pattern's authenticity. The evolution of the members of the warring lineages described by Dawkins produced authentic-patterns, but the four 'great steps' described by Wilson did not. Although we had good reasons to think that Wilson's pattern was largely impressionistic, this was not because its stages could not be conceived as features belonging to a single evolutionary lineage; Wilson's 'stage 4' - organisms possessing a mind- clearly *did* descend from some of those which evolved through the Cambrian explosion ('stage 3'), these were all descendants of the first eukaryotic organisms ('stage 2'), and *all* organisms are descended from the those present at life's inception ('stage 1').

The unique natures of evolutionary lineages means they require special handling, and the awareness that despite the genuine relationships which hold between the organisms and species belonging to a single evolutionary lineage, our ways of presenting and interpreting the events which occur over a lineage's evolution can mislead. As we observed in chapter 3, progress in evolution is both real and widespread: selection advantage and functional improvement often coincide and traits get better at performing their functions as a result. With this conception of evolutionary progress secure, this chapter will be concerned with *procedure*; specifically the procedures used to convey the nature of evolution within lineages and how they can both accurately and mis- represent the progressive nature of the evolution they describe.

Our first lineage-specific procedure is *delineation*: the process of identifying out of the many splits and diversifications that occur to an ancestral population, the narrow path that leads back and links the later descendant population with that of its evolutionary predecessors. As we 'pick out' a part of the natural world when we do this the procedure should be uncontroversial, however as the delineation of a direct ancestral line is often preparatory for considering the identified line as somehow more important or 'central' than those which 'deviate' from it, delineation is often regarded with suspicion. But the *identification* of an evolutionary lineage and the *privileging* of an evolutionary lineage are separate procedures and we should not confuse them. We can recognise the line of descent stretching back from contemporary humans to their ancient ancestors in the distant past, but this is very different from thinking of that lineage in particular as more important, 'central', or 'progressive' over and above the vast number of alternatives available. All extant organisms possess lineages equally as ancient as ours, and many will feature characteristics which are just as unique and - arguably - just as unlikely and complex. Our appraisal of Wilson's four steps showed that if there are reasons for thinking that our lineage does have specific properties which justify its separation from those which comprise the remainder of terrestrial life Wilson does not adequately provide them.

Our second lineage-specific procedure is *demarcation*: once a lineage *has* been identified, there must be good reasons to divide that lineage up - to segment it - into the particular sections and stages chosen. Wilson's chosen lineage bookended by the 'human mind' and the 'advent of life' was segmented twice producing three distinct sections. As we saw, Wilson appeared to lack a consistent rationale for dividing the lineage up in this particular way: employing the first eukaryotic organisms for this purpose was dubious, while omitting the advent of sexual reproduction appeared remiss.

The criteria of *delineation* and *demarcation* are recognisably distinct, and an appreciation of both is critical for a balanced appraisal of the evolution of lineages and the organisms and species which comprise them. If we do prioritise certain lineages over others good reasons must be provided, and if we do segment a particular lineage in a certain way, we must also have good reason for thinking it preferable to the many other ways in which a lineage can be segregated and 'indexed'. But while the motivations behind the *delineation* of certain lineages has received considerable philosophical attention, those behind *demarcation* are frequently overlooked. Wilson and nearly all accounts of evolutionary progress approach the problem by attempting to justify the specific promotion of the human lineage from other evolutionary lineages - for finding reasons to think that human beings are special in some way - but even rare approaches which recognise 'evolutionary advance' less anthropocentrically can mislead when the reasons for their specific demarcation are not given due care and attention. Our aim in this chapter will be to consider progress in evolution over a longer-term perspective by looking at the delineation and demarcation of lineages in turn.

This chapter will proceed as follows: In **5.2** we will identify the *delineation* of evolutionary lineages as a valid operation: lineages are objective parts of the natural world, and treating them as such should not generate the controversy that it often does. We do however need to provide good reasons to prioritise certain lineages over others, and have cause to be particularly weary when the lineage in question is that of humans beings- although this is not sufficient reason not to prioritise the human lineage if alternative justification can be given. In **5.3** we will discuss the second procedure: *demarcation*. As lineages can be

produced out of countless individuals and many species, the objective nature of the ‘stages’ through which a lineage evolves will be less straightforward. Lineages can be demarcated in many different ways and often the reasons for dividing a lineage will be so as to characterise the lineage’s evolution as a ‘progressive’ process; changing the characteristics of the organisms and species earlier in the lineage so that they resemble the characteristics of the organisms and species at the lineage’s close. In 5.4 we will use the findings of the Long Term Evolutionary Experiment (LTEE) to demonstrate that although evolution is not forward-looking, this procedure of demarcation often *is*. When the demarcation of a lineage is done so as to explain the incremental evolution of the lineage’s later stages, the appearance of progressive evolution ‘towards’ the lineage’s later stages will be the result. However, as we will appreciate through the work of Nilsson and Pelger in 5.5 reality and representation can coincide; a lineage can *look* progressive for the reasons stated in section 5.3, but actually *be* progressive for different reasons as stated in chapter 3; later stages of the lineage may be functionally superior to those which proceed them. In 5.6 we will head-off a potential problem with this chapter’s findings: our equating of evolutionary progress with functional improvement, as this reliance on the SE conception of function might appear to expose a significant weakness in our approach: it would seem to prevent the comparisons of the functional performances of distantly-related species, and so severely limit what counts as evolutionary progress as a result. In response we will observe that although the SE conception is sufficient to tell us what the function of a trait is, it does not define what the abilities which are functional are- it is possible to identify these abilities in widely-different species independently of the SE approach. 5.7 Recaps this chapter’s findings and concludes.

5.2 Delineation and anthropocentrism

Over the course of evolution what began as simple and relatively uniform living forms have greatly multiplied in number and diversified in form. The reproduction of even the simplest of organisms is far from perfect and over multiple generations these differences accumulate; a process which is compounded in sexual species by reproductive isolation when differences in genotype cement speciation. Due to the successive splitting of species from a small group

of common ancestors, net bio-diversity has increased over evolutionary time; although of course many lineages have ended and continue to end in extinction. The traditional representation for this spread of diversity is the ‘tree of life’; with life’s ancient ancestors nestling at the base of trunk with phyla, classes, orders and other taxa symbolised as boughs and branches, ending in the twigs and tips of contemporary species. Although widely used, this metaphor has fallen out of favour as those ever-alert to signs of latent progressionism in evolutionary thinking (O’Hara 1992, Gould 1996c, Ruse 2009) object that the positioning of earlier taxa at the base of the tree suggest that later taxa are ‘higher’ or superior in some way to those below. The placing humankind at the top of the tree- its inevitable location - being a particularly egregious example.

Our purpose in this section will not be to defend the ‘tree of life’ representation. The concerns voiced by Gould et al. are not the tree’s only weaknesses³² and in the light of the availability of alternative and non-hierarchical representations³³ we all might agree that the usefulness of the metaphor is recognisably limited. For those lineages featuring organisms whose genomes are too enclosed for horizontal transfer, once a lineage splits, it splits for good and the tree metaphor can work reasonably well in these circumstances. But conceiving of evolution as a series of bifurcations and diversification like this can mislead in other ways; it is often *contrasted* with a linear interpretation, as though the idea of species occupying the tips of the tree having a direct ancestral line leading back down to the branches and limbs into their ancient evolutionary past were somehow problematic.

This claim may seem surprising, perhaps even alarmist, but it is an impression that can be hard to avoid. Cognitive scientists Novick, Shade & Catley (2010, p3) for example state that: “people’s understanding of evolution as a process [which]... specifies that one species undergoes changes that lead it to turn into another species” is a “persistent misconception”; evolutionary biologist J.D. Archibald (2014, p19), praising Gould’s well-known criticisms

³² Those concerned with an accurate depiction of phylogeny at the smallest of scales for example object that the horizontal transfer of genetic information by the ‘bacterial sections’ of the tree makes the notion of ‘branches’ incoherent (McInerney et al. 2008, Puigbo 2009, O’Malley, Martin & Dupré 2010);

³³ Circular representations of the phylogenetic diversification from a last common ancestor have become increasingly prevalent. Here later species are depicted radiating out from a central point, rather than evolving ‘upwards’.
<https://www.sciencemag.org/news/2015/09/first-comprehensive-tree-life-shows-how-related-you-are-millions-species>

of linear ape-to-human illustrations of evolution, comments that “...this type of representation fuels profound misunderstandings” and Gould himself regarded the depiction of evolution as a linear process with consistent hostility: “ladders are false abstractions, made by running a steamroller over a labyrinthine pathway that hops from branch to branch through a phylogenetic bush” (Gould 2002b p97); and that “Each tip can be connected to a last common ancestor by a labyrinthine route, but no paths are straight and all lead back by sidestepping from one event of branching speciation to another, and not decent down a ladder of continuous change” (Gould 1996, p67).

Yet evolutionary lineages are of course quite real, and a recognition of their reality does not require a commitment to the view that species ever evolve in ‘ladders’ *without* successive splitting- although debate on the coherency of this view continues (Futuyma 1987, Hallam 1997 & 2009, Turelli et al. 2001, Emerson & Patino 2018)- only that the lineages of more recent species are traceable back through an unbroken line of inheritance charting their descent through successively more ancestral populations³⁴. Much as the ‘tree of life’ representation of evolutionary history may mislead us into an implicit acceptance of Homo sapiens’ ‘higher’ status, Gould’s legacy of concern here appears to be that conceiving of evolution as producing linear series may similarly mislead us into thinking of evolution as a directional force - progressively driving species to ever-greater perfection. In what has remained one of the clearest and best-known attacks on the position that biological evolution could be characterised as a ‘progressive’ ladder of improvement like this, Gould declared progress to be:

“...a noxious, untestable, non-operational, intractable idea that must be replaced if we wish to understand the patterns of history” (Gould, 1988b. p319)

³⁴ The term ‘lineage’ will be employed in this dissertation in an general sense and refer to any ancestor-descendant relationship between biological entities at the same ‘hierarchical’ level. Lineages may be comprised from ancestor-descendant sequences of species, populations, organisms or traits; the work ahead will utilise all of these. Typically lineages of genes would also be included into such an expansive account but will be omitted here as our discussion centres on the nature of *changes* occurring to the members of a lineage, rather than their endurance. Unlike the members of the lineages of species, populations, organisms or traits, gene lineages by convention are comprised from members that have *not* changed.

And we are left in little doubt as to what Gould's somewhat reactionary approach is a reaction to; the psychological target motivating this attack follows shortly after:

“Our geological confinement to a moment at the very end of recorded time must engender suspicions that we are lucky accident, an afterthought rather than the goal of all creation. Progress is the doctrine that dispels this chilling thought- for if life moves inexorably forward, however fitfully, towards the ultimate embodiment in human consciousness, then the restriction of homo sapiens to a final moment poses no challenge to the general hope; for all that came before may now be interpreted as part of a process scheduled to yield our form from the start.” (Gould 1988b p319).

When we consider the ‘great steps’ of Wilson - culminating in the ‘origin of the human mind’ (Wilson 2001); the stages of Maynard-Smith and Szathmáry’s ‘Major Transitions In Evolution’ - culminating in ‘human language and society’ (Maynard-Smith and Szathmáry 1995); and J.Huxley’s claim that progress in evolution consists in ‘greater control of and independence from the environment’ - and which human beings have achieved to the greatest extent - (J.Huxley 1942); it is difficult to deny that Gould may have a point, and that when evolutionary biologists do stray into more-philosophical waters, the exceptionalism provided to human beings and characteristically-human traits may not be entirely unbiased.

As our findings in chapter 4 demonstrated, our dispositions as human observers can easily mislead, and from this perspective Gould’s position and ours have much in common. For us, the psychological bias for pattern recognition may lead us to erroneously think of entirely phenomenal patterns as authentic, for Gould a kind of species-wide disposition to ‘existential angst’ biases our appraisal of evolution to regard Homo sapiens as special in some way whereas in fact we are no more and no less a product of capricious evolutionary forces than any other species. In both cases human biases threaten to distort an accurate understanding of empirical reality. As we saw in chapter 3 however an evaluative approach to the products of evolution should not be ruled out preemptively. The traits of some organisms perform their functions better than the traits of some of their conspecifics;

adhering to the ‘selected effects’ conception of biological function requires an implicit acceptance of this. But notice that holding that certain *traits* are functionally superior to other traits of the same type is a very long way from holding that certain *species* are superior to other species, as per Gould’s criticism. As the SE conception identifies the functions of traits with their selectively-beneficial historical effects it would actually seem to make cross-species comparisons like this highly problematic; given that different species- and therefore the functions of the traits of different species- will have distinct selection histories. If we were to carry out cross-species comparisons in the manner that Gould is objecting to, it seems we would need to provide alternative, or at least supplementary means, for their justification as the selected effects approach will not achieve this unaided. We tackle this problem in 5.6.

Contra-Gould however, there are two further points to make here:

Firstly, notice that Gould is in danger of conflating a justifiable recognition of our disposition to regard phenomena in a certain way with an automatic dismissal of the reality of the phenomena itself. It could well be the case that human beings are psychologically inclined to count themselves as distinct and ‘superior’ in some way to the rest of the natural world *and* for this belief to be warranted. The various reasons given by the above authors for humanity’s special position in their various classifications certainly seem to have *some* initial plausibility and not to be ‘spun out of whole cloth’. Huxley’s contention that evolutionary advance consists in an increasing control of and independence from the environment for example, appears to be an empirically-measurable quality - rather than entirely impressionistic one as Gould’s pejorative ‘non-operational’ comment might suggest - and one which would be selectively-beneficial over a range of environments given that the detrimental effects of a capricious environment could be mitigated to some degree. And despite Maynard-Smith and Szathmáry’s established problems with consistency (Calcott & Sterelny (eds.) 2011, Szathmáry 2015, O’Malley & Powell 2016), it seems clear that out of the three properties which they see as having increased over terrestrial evolution: the loss of independence of replication in the units of higher collectives; an increase in division of labour; and an increase in the efficiency and storage of information; all three are potentially

measurable, and the final two seem specifically to be characteristics which human societies are especially proficient at. Clearly these authors need to provide subsidiary reasons why *these* properties in particular should be used as a standard against which the whole of the living world may be measured as opposed to other properties where Homo sapiens may fare more poorly. Although we might suspect that this cannot be done impartially, it would be wrong to dismiss the possibility out of hand.

Secondly, although it is common to conflate the possibility of anthropocentric bias with the perception of evolutionary progress as does Gould, this need not necessarily be the case. As Darwin himself observed:

“To attempt to compare members of distinct types in the scale of highness seems hopeless; who will decide whether a cuttle-fish be higher than a bee?... In the complex struggle for life it is quite credible that crustaceans, not very high in their own class, might beat cephalopods, the highest molluscs; and such crustaceans, though not highly developed, would stand very high in the scale of invertebrate animals”

(Darwin, p308, 1872 {6th edition}).

The urge to rank different species hierarchically is quite general, and not just an enterprise in seeking-out features which humanity has in abundance and which makes them allegedly superior. The desire to rank the natural world still occurs when Homo sapiens is *not* one of the species whose advancement is being adjudicated. Just as the above authors need to provide empirically-justifiable reasons for placing humanity at the latter end of the evolutionary scale, those wishing to compare the ‘advancement’ of other non-human species must also provide similar reasons for their relative positions. Although as recognised by Darwin, the hope that this procedure could be carried out in a manner which is scientifically respectable seems remote.

So the intuitions of those who feel that life has ‘advanced’ over evolutionary time should not be conflated with impossibility that life actually has advanced in some way; nor is the practice of hierarchically-ordering the different species of life necessarily motivated by a

desire to confirm human superiority: although this may often be the case. Additionally, although Wilson et al. have no solid justification for *prioritising* the specific lineage stretching back from Homo sapiens to the beginning of terrestrial life over and above the vast number of alternatives available, this should not blind us to the fact that this lineage can be *identified*. It is not the identification of the human lineage, or the identification of evolutionary lineages *per se* which is problematic, but the danger that focusing on certain lineages in particular will mislead us into conceiving of some lineages as more *central* to the history of life than others. This of course is false, and Gould is right to criticise it. But neither is Gould's characterisation of evolution as a 'labyrinthine route' quite right either. There is a definitive and real direct ancestral line connecting modern species to their ancient predecessors and it is equally as wrong to suggest that the procedure of 'delineating' these evolutionary sequences from those of their phylogenetic neighbours should be considered invalid or misleading. As long as the relationships between the species and organisms under examination genuinely are ancestral this is sufficient qualification for the lineage connecting them being an objective and real part of the natural world rather than imposed and artificial as Gould's criticisms seem to imply.³⁵

5.3 Demarcation: An A to Z of linear evolution.

As a lineage evolves, species and traits come and go, and this successive replacement of traits and species permits a conception of evolution as a kind of *directional* process, acting to move from one successive form on to the next. It is clear how the gradual replacement of ancestral types by their descendants might be preemptively construed as 'progressive': as those later in the lineage *replace* those earlier, it may seem natural to think that this was achieved by virtue them being *superior* to their predecessors in some way. Along with anthropocentrism, this somewhat lay interpretation of evolutionary history has acted as one of the premier targets towards which opponents of the 'progressive view' of evolution have

³⁵ How we set about the process of delineation will depend on what we intend the resultant lineage to be a lineage of; as will the number of identifiable lineages produced. Clearly as species are comprised from populations, which are comprised from organisms, which themselves possess traits, there will be far more lineages of traits and organisms than there are lineages of populations, and more lineages of populations than there are lineages of species.

directed their attention (Ayala 1974, Jablonski 1986, Gould 1996), and one which even those generally more sympathetic towards an evaluative interpretation of evolutionary phenomena recognise as a gross oversimplification (Dawkins & Krebs 1979, Ruse 1993, Dennett 1995a, Dawkins 1997). Descendant organisms succeed their predecessors for any number of reasons: sheer chance, geographic distribution, changes in climate and in ecological environment and so on, and much anagenetic evolution may have no bearing at all on any systematic differences in fitness or phenotype between earlier and later lineage members. Gould and those influenced by him are again quite right here; holding that later members of a lineage must be 'superior' in some way by virtue of their succession alone is not only philosophically problematic, but in a large number of cases will be straightforwardly incorrect. But equally, in some cases, later lineage members will succeed their predecessors because they *are* better able to survive and reproduce. The replacement of earlier types by their descendants will not always be due to haphazard factors unrelated to fitness such as the effects of a capricious environment. Later lineage members may well be 'superior' in some way to their ancestors by virtue of being able to perform certain functions that their predecessors lacked, or being able to perform certain functions to a higher standard. But divorcing genuine cases of 'evolutionary advance' like this, and cases which look very like evolutionary advance but are not is often far from straightforward.

Any lineage which features significant evolutionary changes will be comprised from a vast amount of individual organisms, permitting numerous ways in the which the changes which occur as the lineage evolves may be emphasised and interpreted. Even when the stages 'bookending' a particular section of an evolutionary lineage can be empirically established, the nature of intervening events can be highly speculative and controversial. Consider the 'Aquatic Ape' thesis for example: no-one credibly disputes that the lineage culminating in *Homo sapiens* stretches back to ancestral apes, but the suggestion that this involved an intermediate 'aquatic stage' is contentious; there is currently *not enough* corroborative evidence for it (Morgan 1982, Langdon 1997, Bender et al. 2012). This can be understood straightforwardly as empirical underdetermination; if further evidence were to appear we might use it to strengthen or weaken the conjecture about what occurred over the lineage's intermediate states. But there is a more interesting problem that arises when an evolutionary

lineage provides us with *too much* information for any one investigation to comprehensively cover. In such cases, what we can characterise as an ‘editorial decision’ must be made: what to omit, what to include, what to emphasise and what to de-emphasise, and this procedure presents separate, philosophical difficulties. Editorial decisions require editors to make them and what counts as an appropriate ‘edit’ of a given lineage for one project may be entirely inappropriate for another with different explanatory aims. There is no objective and perspective-free way to go about this. In contrast to an underdetermined conjecture like the ‘Aquatic Ape’ here we proceed taking evidence *out*; and when we want to break down a large and continuous evolutionary lineage into its component stages the relevance of the particular stages chosen will always be project specific.

For example: suppose that we miraculously had access to a complete fossil record spanning from the basal whale lineage member: Pakicetus, to the first of its whale-like descendants, a collection comprised from an enormous amount of individual specimens. In order to employ the collection for any explanatory purpose (other than showing the sheer size of the lineage), we would first have to edit the specimens down to a much more manageable number. How we would go about this procedure would depend on what *sort* of information we are hoping the chosen specimens will provide. We may feel it representative to pick samples separated by one million years, or perhaps by one million generations, and these might tell us much if our interests concerned the rate of aquatic adaptation, or the order in which each adaptation occurred. But for the majority of evolutionary explanations this approach would be inadequate. What if we were interested in specific adaptations for life in the oceans such as the disappearance of external ears and hind limbs, the appearance of tail flukes and the bulk of these developments fell between our one million markers? Any explanation for how the lineage adapted to an aquatic environment would be severely compromised without including such events so our decision to demarcate the series using a metric blind to them would be procedurally inappropriate. Balanced ‘curation’ is required: a set which is large and detailed enough so as to represent all the significant developments, but small enough so that the number of entries is limited to a practical and informative amount.

It might seem natural in these cases to use the ‘key stages’ which occurred to the lineage as it gradually adapted to a fully aquatic lifestyle as a guide for demarcation; indeed when the changes undergone by an evolutionary lineage are quite radical this is typically how it is done (Jenner 2018). But notice how different ‘insensitive’ scales of measurement like time and number of generations are - which provide a fixed scale against which the lineage’s evolution can be objectively measured - from ‘sensitive’ ones which take into account qualitative features of the evolving lineage. While the former have some flexibility in that duration and number of generations can be ‘set’ to provide the grain of detail required, the resulting edit will be blind to any evolutionary changes the lineage has undergone. This is fertile ground for potential conflict: *independence* of measurement is of course something to be strived for in many cases - empirical science would be impossible without it - but employing a systematic means of measurement for unsystematic phenomena like evolutionary change will often result in an ‘edit’ that is inadequately *representative* of the phenomena we are trying to understand.

We can illustrate this as follows. Let us suppose that the Pakicetus-to-Whale lineage can be comprehensively represented by twenty six stages; starting from specimen A at the start of the lineage, through the letters of alphabet to specimen Z at the lineage’s close³⁶. A systematic demarcation of this lineage by fixed markers- years or generations- would result in the following kind of structure:

Edit 1: **A** B C D **E** F G H I **J** K L M N **O** P Q R S **T** U V W X **Y** Z

We place a marker (in bold) at every 5th sample to produce ‘Edit 1’: A-E-J-O-T-Y. We might wish to get more detail from this lineage and set our demarcation at every 4th letter, or feel that whatever task we are using the lineage for can be just as effectively carried out by selecting those at every 6th or 7th. The point to bear in mind here is that whatever the

³⁶ We are aware of course that in actuality *this* A to Z representation in common with the ones we are describing can only be reduced from many millions of samples according to some preference as to why those twenty six are chosen in the first place. I feel the illustrative benefits of this example outweigh its slightly self-referential nature.

chosen grain of our measurement it will provide for us an objective result: evolutionary change in the lineage occurring in the gaps between the chosen markers.

In contrast, an investigation into aquatic adaptation within the same lineage using the evolutionary developments themselves to guide its demarcation not only provides a different result, but provides a different *kind* of result. This alternative approach to the procedure might produce something like this:

Edit 2: **A B C D E F G H I J K L M N O P Q R S T U V W X Y Z**

A: Basal Pakicetus specimen

H: Limbs much reduced in size

I: Bone thickens around ear to aid hearing underwater

P: Nasal opening has shifted upwards and backwards along top of the head

T: Eyes of sample have moved to the sides of the head

U: Tail is now recognisably fluked

X: Hind legs completely disappear

Z: Last common ancestor of all whales.

In 'Edit 2' the lineage's gradual evolution away from a terrestrial environment 'towards' an aquatic one is charted throughout each of the eight selected stages. Sample H (marking the reduction in limb size), and sample P (marking the relocation of the nasal opening) for example are evidently of major importance for this particular project, as are U and X, but they are also ones that would have 'slipped through the net' had the lineage been demarcated systematically as it was in Edit 1. Edit 1 is demarcated according to time or number of generations, whereas Edit 2 is demarcated according to the features of the evolving lineage itself. This contrast can be more clearly appreciated by considering that if the evolution of the lineage in question had been markedly different, Edit 1 would produce different results - the organisms and species which comprise the lineage would after all be different organisms and species - but the *demarcation* of the lineage (A-E-J-O-T-Y) would

remain the same, this is because the procedure that produced Edit 1 did not take the evolution of the lineage into account. This would not be the case for Edit 2; the differences in evolutionary development would not only produce a different result as in the revised Edit 1, they would *additionally* make Edit 2's original demarcation (A-H-I-P-T-U-X-Z) defunct. This is because the procedure that produced Edit 2 does take the evolution of the lineage into account. In Edit 2 as the demarcation is carried out according to the evolving characteristics of the lineage, when these change so does the edit. Contingent upon how major the evolutionary differences in the revised lineage are some of the stages originally described in Edit 2 might not even exist.

Although the 'impartially' of Edit 1's demarcation may strike us as more scientifically respectable, there is no one 'right' way to go about the procedure. Edit 1 and Edit 2 produce different results because they are intended to measure and describe different things; for Edit 1 this is the rate at which evolution has occurred; while for Edit 2 it is *how* the lineage transferred from a land based- to sea based- lifestyle that is of primary interest. If it is evolutionary rate we wish to measure, the approach of Edit 1 will be the correct one to use; alternatively, if it is qualitative features of the evolving lineage itself that we are interested in, the approach of Edit 2 will be appropriate. An edit tasked with concisely presenting a lineage's evolutionary transformation is obliged to demarcate using unevenly spaced and grouped markers like in Edit 2 because this is how the actual evolution of the lineage is very likely to have occurred; in unpredictable and sporadic bursts and pauses. The evolutionary 'stages' undergone by a lineage featuring significant change will almost certainly not accord with anything like the neatly-spaced temporal markers as described by Edit 1, we need a bespoke edit of samples chosen with these related changes *explicitly* in mind if we are to present this lineage's transformation satisfactorily.

Now, we might have concerns that by rejecting an impartial metric the representation of the evolution in question will be unavoidably arbitrary and subjective but this is not exactly right. Whether we choose to edit an evolutionary lineage so as to accord with certain changes that it has undergone is up to us - but once a particular explanatory project has been decided upon, it seems that what counts as significant *within* that particular project is not

something that we are similarly free to opt in or opt out of. For example: Our ‘comprehensive’ lineage of x26 samples had been substantially reduced to x8 in Edit 2 because it was those samples rather than any of the excluded eighteen that ‘fitted’ with our project to understand how it was that a largely land-based organism at stage ‘A’ evolved to become one fully adapted to an aquatic environment like that at stage ‘Z’. There may of course be some disagreement about the relevance of specific samples in any given explanatory project, but these will not simply be due to the whims of the investigator, but because of disagreement over how much they add to an understanding of the lineage’s transformation over and above that of the samples whose inclusion everyone can agree upon. Expanding the edit to include non-relevant samples may compromise the project by introducing unnecessary ‘noise’, making the evolutionary sequence of interest less apparent. To use another example: say the sample ‘D’ represents the first in the lineage to reach 2.5 metres in length from nose to tail. As the A to Z lineage is an objective ordering, adding D into our edit of aquatic adaptation to make A-D-H-I-P-T-U-X-Z is not *technically* inaccurate as it is justified by the lineage’s phylogeny, but D does not ‘fit’ in that same position as comfortably when the lineage is considered *thematically*. A 2.5 metre body length is of course just one of many properties that specimen D possesses, but if none of D’s properties add to our understanding of aquatic adaptation over and above those of the original eight specimens, the act of including D in the sequence would seem gratuitous; we would be fully justified in asking what purpose the inclusion of D in the re-edited sequence performed.

Equally, our demarcation of the lineage according to aquatic adaptation would also be compromised if we were to *leave out* certain specimens. Suppose that we now take Edit 2 and remove specimen X (the disappearance of the hind legs), leaving the reduced sequence A-H-I-P-T-U-Z. We might still have a sequence recognisably depicting the lineage’s aquatic adaptation, and perhaps if we were limited to seven stages the most informative one that could be produced, but it is evidently not the most informative one overall because a major event in the sequence is missing. Looking over the shortened edit we would be entitled to ask why the final disappearance of the hind limbs had been left out. That A had legs so as to move around its terrestrial environment, and Z did not is a major component of the transformation that the edit was set up to show. The purpose of editing the larger collection

was to explain how this, along with many other land-to-sea transformations took place, and once this objective has been decided on, the inclusion of certain specimens like X seems to be near-mandatory. An edit which omitted the disappearance of hind limbs, like one that omitted the addition of the fluked tail or the relocation of the nasal opening, would be severely- and *obviously*- incomplete as it would fail to accurately represent the evolutionary development it was explicitly set up to illustrate. Once we have locked onto an explanatory target we have a reasonably good understanding which out of a give range of specimens will be appropriate for it. If this is ‘aquatic adaption’ then this means a definite ‘yes’ for the inclusion of hind-limb disappearance (X), and a definite ‘no’ for the first specimen to reach 2.5 metres (D). So any concerns that the demarcation of a lineage according to qualitative features of its evolutionary development will produce *wholly* subjective results should be mitigated to some degree; we cannot simply decide to omit or include any stages within an edit arbitrarily.

Freedom of choice appears to be only unconstrained when it comes to the delineation of a given lineage and which of its ‘evolutionary themes’ we find worthy of further investigation, but this is not as straightforward as it may seem either. Notice what motivates us to embark on the project of editing down an evolutionary lineage in the first place: we want to understand how various evolutionary influences combined to change the characteristics of the species and organisms earlier in the series to produce the distinct characteristics of the descendant species and organisms which succeeded them. To return to our A-to-Z example: in common with all living things the whale has an evolutionary history, we want to understand how a creature of that size and with those traits and that physiology and so on could have evolved from a creature very unlike it in many ways by a series of incremental steps. It is the development of the properties possessed by the *later* samples of the series when contrasted with the properties of the earlier specimens that motivates us here, and it is the differences between the specimens at either ends of the sequence and how they were bridged that the procedure of editing-down an extended lineage is particularly good at illustrating. In the case of the whale, our chosen ‘theme’ of aquatic adaptation seems to be quite a suitable one, as it covers many of the differences between the properties of the modern whale when contrasted with those of Pakicetus at the start of the sequence. But this

is not the only interpretation for the changes that the whale lineage has undergone. 'Size increase' could provide an alternative - although less compelling - theme for the linking of the lineage's two 'extremes' as the intervening stages show how Pakicetus grew incrementally over evolutionary time to the size of the whale's last common ancestor; we see how the contrasting sizes of the last and first specimen are bridged by their intermediaries showing size increase over the span of that section of the lineage. Although any given lineage, and especially one as lengthy and complex as the transformation of Pakicetus-to-whale, will have a variety of different 'thematic' threads weaving through it, some themes will do this more more straightforwardly than others. For our A-to-Z lineage the clear front runner is 'adaptation to an aquatic environment', with perhaps size increase as a distant second. Other than in lineages which are very brief, or those where the changes between earlier and later lineage members are very limited, there will be multiple 'thematic' ways of interpreting the evolution which has occurred, and so multiple ways of editing the lineage so that these various themes can be highlighted and represented.

So a lineage is edited down to a more manageable size by its division into relevant stages, and the relevance of any stage with a given edit will be in accordance with the lineage's overarching theme- many of which might be suitable. But notice also the side-effect that viewing the evolution of a lineage through a sequence of thematically-edited stages has on the way that we conceive of the actual evolution of the lineage itself. Because the editing is done so as to thematically link the characteristics of the organisms and species later in the series with those of their predecessors, the reasons for the inclusion of earlier samples into a given edit- their 'thematic relevance' - is property which is only attained by those particular stages some time *after* the evolution of the stages themselves.

As is often stressed evolution does not plan ahead, but the seldom-recognised corollary of this is that the significance of evolutionary events is something which can only be appreciated retrospectively. Our understanding of Z's gradual evolution for example is informed by events which occurred temporarily before it, but we are able to know *which* events will prove to be important for this later stage as we can view the lineage's evolution from a privileged position: we know that they lead to Z. But what we identify as a

significant event to explain the evolved properties of Z or those of any other subsequent state will not have that relevance *at the time of their occurrence*; the events occur first, then, sometimes, subsequent developments may lead us to reinterpret their importance.

As has been widely appreciated (Beatty 1995, 2008 & 2016, Cleland 2011, Powell 2012, Currie 2014, Hesketh 2016, Blount et al. 2018) there are clear parallels with the way in which we retrospectively interpret the evolution of complex and unlikely traits and adaptations and the way in which we reinterpret historical events through the perspective of their later consequences. We are now able to recognise Caesar's crossing of the Rubicon for example as an event of vital importance in the ending of the Roman Republic and the establishment of the Roman Empire but this was far from clear at the time of its occurrence. The first Roman emperor was finally given the 'power to command' by the senate at the end of a protracted civil war, 22 years after the fact; only *then* did the importance of Caesar's transgression become recognised. With these parallels in mind it should be unsurprising that when a lineage is edited down in accordance with an overarching theme, the resulting edit can make the lineage's subsequent evolution appear predestined or inevitable. When viewed retrospectively, the intermediate stages which describe the transformation of an evolutionary lineage much like the intermediate stages which build up to historical events can appear to be 'progressive' because in the sense of being future-orientated the mode of explanation guarantees that they genuinely *are*. The inclusion of earlier, antecedent stages in both this type of evolutionary explanation and historical explanations are done explicitly on the basis that they help explain how events later in the sequence came about.

The fact that the processes underlying the evolution of lineages have no foresight, but that our explanations for the evolution of lineages *rely* on 'future knowledge' in this way creates something of a disparity between the properties which the phenomena appear to have by virtue of our means of explaining them and the properties of the phenomena itself. It may look, for example, as if the incremental evolution of the whale through the diverse forms of its predecessors was 'progressive' in that the stages preceding 'Z' gradually build towards the lineage's final form. Our explanation operates on this basis, but as we know, evolution itself does not. We are able to interpret the intermediate stages P, U, X etc as 'building'

towards the lineages' final stages because we possess the foresight that the process itself lacks; P, U and X did not evolve *because* they were stages 'on the way' to becoming a whale, although our mode of explanation for the lineage's evolution treats them as though they were.

It might be tempting here to think that the problem of progress in evolution is not one big problem as is commonly construed but is better-characterised as many smaller and related problems; most of which we have either answered or dismissed. We saw earlier that Wilson's approach was ill-defined and the patterns he described lacked a higher-order mechanism; we saw that Dawkins's defence of evolutionary progress aligned closely with our justification of normative language as both recognised that to be cumulative, adaptation must not only provide a selective benefit, but provide a selective benefit through functional improvement; we saw that Gould's characterisation of progress as a kind of existential palliative was neither sufficient to say that progress does not exist, nor was it necessary as non-anthropocentric perspectives on progress date back to Darwin if not before; and now we have seen that the lay conception of progress as witnessed through a lineage's stage by stage progression into later forms can be artifactual, and due more to our methods of explanation, than to the properties of the evolutionary process itself.

Once these observations have been taken into account we may think that there is little of the concept of evolutionary progress left to examine, but this would be to re-make Gould's mistake. We know that the process of demarcation divides the lineage into stages that we retrospectively characterise as progressive in the light of their future effects, but to dismiss the evolution of *all* lineages as non-progressive because they can be characterised in this way mirrors closely Gould's claim that because humans may think themselves superior to the rest of the living world, this is enough to dismiss the idea out of hand. It could still be the case here - as it was there - that the evolution of a lineage can be characterised as progressive due to our explanatory approach *and* for the evolution which the lineage has undergone to be actually progressive in fact. The former does not rule the latter out.

Although the gradual adaptation of the whale lineage to an aquatic environment certainly seems progressive when it is demarcated into stages designed to highlight that very characteristic, the question here should be: do we have reasons to think that the underlying evolution is progressive *in addition* to those produced through the means of our explanation? Due to our findings in chapter 3 we have some guidance as to how we might go about finding this out. When a lineage is demarcated in order to show how its later stages were produced through the evolution of those earlier, do these stages also chart an *improvement* of some sort? Some lineages may feature the evolution of traits which are better-functioning than those which preceded them and so feature improvement in this limited sense by doing so, but there are major difficulties in transposing this relatively 'localised' type of functional improvement onto lineages which extend over much longer periods of time and whose earlier and later members are very different. As we saw earlier, our grounds for thinking that a comparison of the performances of traits could be carried out naturalistically relied heavily on the conception of functions as 'selected effects', but as ancestral traits become separated by ever greater amounts of time and shaped according to different selection pressures these grounds become very much reduced.

These difficulties are particularly obvious when we attempt to compare the performances of the traits belonging to species as different as Pakicetus and its whale-like descendants. For many traits no comparison will be possible as they do not belong to both organisms - Pakicetus did not have flippers and whales do not have legs - but even when traits are retained throughout the lineage, the effects for which they were selected as the lineage evolved may become distinct past the point where a comparison of their abilities would be possible. In contrast to the major overhaul of much of its anatomy over evolutionary time, many of the major organs possessed by Pakicetus: heart, lungs, liver, kidneys and so on persisted throughout the lineage and were inherited, albeit in modified form, by its whale-like descendants. But the *effects* for which Pakicetus' lungs provided a fitness advantage and the effects for which whale lungs provided a fitness advantage will almost certainly be significantly different. What would count as efficient or superior functioning for the lungs of one of these species would almost certainly not count as such for the other - it would very likely prove fatal. Whale lungs are certainly *different* lungs to those of Pakicetus, but they

are not *better* lungs. Our method of analysis developed in chapter 3 relies on the selected effects conception of function establishing that the functions of closely-related traits are comparable, it was not designed for cross-species comparisons like this, even when the species in question are ancestral as there are here.

So when a lineage is as long as that of the whale and the changes through which it has undergone as substantial, even a comparison of the traits which do persist throughout the lineage's evolution will be problematic. This should not surprise us as the impossibility of comparing profoundly different species in different selective environments was made clear by Darwin above. Although cuttle-fish did not evolve into Crustacea whereas Pakicetus did evolve over time into its whale-like descendants, some of the differences between these organisms will be just as pronounced and Darwin's scepticism just as applicable. To know if a lineage's evolution is 'progressive', we must first be able to validly compare earlier and later members, and to do that we must make sure our comparisons are comparisons between instances of the same type of thing. This rule is quite general: the more different items for comparison are, the less validly a comparison of them can be carried out; and this will be the case for different species whether they are direct descendants like Pakicetus and the whale or are not like cuttlefish and Crustacea.

To conclude this section: while the *delineation* of an evolutionary lineage from its surrounding phylogenetic diversity is a straightforward matter of tracing later organisms and species back along their direct lines of ancestry to those which precede them - a procedure which tracks and picks out an objective part of the natural world; the *demarcation* of a given lineage into component stages is more complicated and the objective nature of the 'edits' produced less secure. Lineages can be variously demarcated according to different explanatory aims, but when they are designed to explain how the stages later in the lineage came about they present the evolution featured as progressive - progressing towards the lineage's later stages. Demonstrating any 'progress' brought about over and above that automatically provided by the form of explanation can be problematic especially when the differences between ancestral species and their descendants are particularly pronounced. As we will see in 5.6 however, these difficulties are not insurmountable in all cases. Sometimes

the performances of the traits of earlier and later lineage members *can* be validly compared, and when the performances of the traits improve, the evolution which the traits have undergone is firmly progressive.

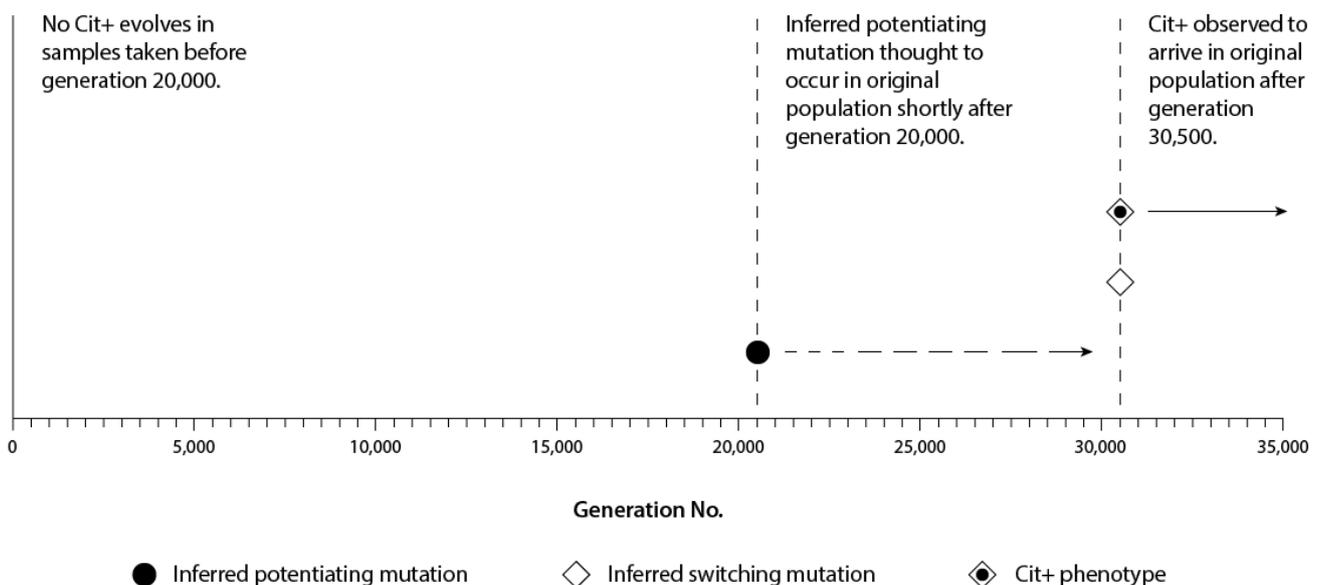
But first we will witness the procedures of delineation and demarcation in action, and advance from our idealised Pakicetus-to-whale lineage in directing our attention to the ‘edits’ of two genuine lineages pulled from the evolutionary literature: the relatively brief and artificially-contained evolution of the Cit⁺ phenotype in the Long Term Evolutionary Experiment (LTEE), and the much lengthier evolution ‘in the wild’ of the vertebrate eye as interpreted by Nilsson & Pelger. We will see how the procedures of delineation and demarcation have been applied, how lineages appear to us as progressive through the form of their presentation, and ask whether the evolution which the lineage has undergone is genuinely progressive in fact. We will then, in section 5.6, return to the discussion of how cross-species comparisons can be validly made.

5.4 The evolution of Cit⁺

The Long Term Evolutionary Experiment (LTEE) commenced in 1988 with the aim of investigating parallel evolution in twelve clonal lineages of *E. coli* bacteria (Blount, Borland & Lenski 2008, Quadt et al. 2014, Blount 2016). After 15 years and approximately 33,000 generations it became apparent that one lineage ‘Ara-3’ had evolved a novel ability; it was able to metabolise the citrate solution used as a neutral medium in the experimental set up. In common with the other eleven lineages, as samples of Ara-3 were retained and frozen every 500 generations it was possible to retrieve a member of the lineage from before ‘Cit⁺’ had evolved to see if ancestral versions of the strain would go on to develop the citrate-metabolising phenotype in parallel as did its decedents. It was found that out of 72 ‘replays’ ranging from Ara-3’s founding population to generation 32,500, Cit⁺ reappeared a mere 4 times, and in each case this occurred in a population taken from the narrow range between generations 30,500 and 32,500; very close to when the original lineage evolved the phenotype. Based on these results, it was hypothesised that Cit⁺ was only able to occur in

later populations due to the trait requiring multiple mutations; a earlier ‘potentiating’ mutation which accumulated in the population and then a subsequent ‘switching’ mutation which were able to act in concert to develop the Cit+ phenotype. This conjecture was strengthened when further inventive experimentation (this time effectively running millions of replays) showed that while the later the generation, the higher the likelihood that it would develop Cit+, the adaptation occurred in no revived samples at all taken from before generation 20,000; the experimenters understood this to demonstrate that the necessary ‘potentiating’ mutation occurred sometime in 1997, after generation 20,000, and that this apparently unique occurrence was itself a very unlikely event.

Fig 5.1³⁷



Unlike our idealised whale lineage, this real-world experiment allows us to be quite precise concerning the lineage’s properties. The line of inheritance tracing back from the later

³⁷ Fig 5.1: (original to this dissertation). Markings along the bottom axis represent x500 generations; the places in the original lineage where samples of the population were taken and frozen. The broken horizontal line indicates increasing probability of occurrence as the number of generations increases. The solid horizontal line represents actual occurrence in the original Ara-3 population. (Graphic by Tim Ashton, original to this dissertation).

samples to their direct ancestors does not have to be idealised and imagined; as the LTEE was explicitly set up to compare and contrast the differences in the evolution of separate lineages, ensuring that Ara-3 was delineated in fact was a necessary requirement of the experiment.

The Ara-3 lineage also comes to us pre-demarcated: it begins with its founding clonal population, terminates (for our purposes at least) with the advent of Cit⁺ shortly after the sample taken at generation 33,000, and is systematically divided by the samples taken and frozen every 500 generations. In line with our observations concerning the need to keep the potentially vast amount of samples in any given lineage to a manageable amount we can assume that one taken every 500 generations to be a frequency balanced between informational pay-off and practical concerns like sampling time, storage requirements and so on. In theory we would be able to present these x66+ samples in the same manner as we did with the x26 of the A-to-Z samples of the whale's demarcated lineage, and then further edit them down; either systematically like in Edit 1, or according to qualitative features of the lineage's evolution as in Edit 2.

The potentiating mutation suggested to have occurred after generation 20,000 is clearly what we might term a 'key stage' for the later evolution of the Cit⁺ phenotype as citrate metabolism does not appear evolvable without it; but notice that even if the event in question were somehow observed when it occurred shortly after generation 20,000 - like Caesar's crossing of the rubicon - its later significance could not have been identified at the time of its occurrence. Only with the advent of the Cit⁺ phenotype some 13,000+ generations (and x26 samples) later did 'potentiating' become an accurate description with the experimenters able to infer the mutation's existence and the approximate generation of origin. The LTEE's investigation into the evolution of Cit⁺ has to be 'historical' and retrospective like this, because the future significance of the mutations occurring in the Ara-3 lineage was not evident at the time of their occurrence. Notice that the inability of the experimenters to recognise the potentiating mutation's importance when it originally occurred was not due to any epistemological or observational failure on their behalf. Even if the mutation had been witnessed and comprehensively analysed at the time of its

introduction into the Ara-3 lineage its future effects could not have been known as these were far from guaranteed. For Cit⁺ to occur the initial potentiating mutation had to be followed by the ‘switching’ mutation; but the arrival of this second mutation, and therefore the arrival of Cit⁺ itself, was not something that could have been known in advance.

Consider how our observations in the previous section might apply to the evolution at the LTEE. We can see that both the delineation and demarcation of the lineage were carried out as part of the experiment’s procedure. The experiment’s original clonal population was delineated into twelve separate lineages, one of which was Ara-3; then as each lineage independently evolved in isolation they were demarcated according to x500 generational gaps, a systematic measurement ‘insensitive’ to any qualitative evolutionary changes that may occur. When it became apparent that the Ara-3 lineage in particular had evolved a novel trait, the experimenters set about editing down the much larger set of demarcated stages into a much smaller amount: those events which together explain how the Ara-3 lineage evolved over time from an initial population which could not metabolise citrate into later population which could. In keeping with our former observations, notice that due to the retrospective nature of evolutionary explanation the experimenters start with the effect - the arrival of the Cit⁺ phenotype - *then* attempt to identify its antecedent causes: here the switching mutation, preceded by the potentiating mutation.

The relative simplicity of the changes which the lineage required to achieve this ability means the most concise edit possible, that composed from Cit⁺’s key stages only, could be provided by editing the lineage down into three phases: Original population - population with potentiating mutation - population with switching mutation (see *fig 5.1*). Notice that the sole motivation for editing the lineage down according to these stages is due to their ability to interpret the various processes and interactions which occurred in the population over 33,000 or so generations in a concise and convenient way. These stages occurred in fact, so the edit is not ‘wrong’ but it is not the only possible edit and the Ara-3 lineage would have been divided up quite differently if the aims of the experimenters had also been different. The arrival of the Cit⁺ phenotype will not be the *only* evolution the Ara-3 has undergone during its first 33,000 generations, and had an explanation been required to

explain the representation of other genes in the bacterial population, alternative means would have had to have been employed and Cit+'s 'key stages' very likely omitted altogether. As with our interpretation of the whale's evolution according to the overarching 'theme' of aquatic adaptation, here Ara-3's evolution is interpreted according to the 'theme' of citrate metabolism; earlier events in the lineage's history are interpreted as progressive steps towards the production of this trait.

Although considerably shorter and less convoluted than the multiple adaptations of the whale lineage Ara-3's evolution certainly *looks* progressive when edited into the key stages for the Cit+ phenotype. But is the lineage's evolution *actually* progressive in fact, or is this appearance due to the method of explanation alone? There is no straightforward answer to this. As the Ara-3 lineage is comparatively brief, we can be confident that its later and earlier populations can be correctly compared as the differences between them will be considerably less than those of the earlier and later members of the whale lineage which made their comparison problematic. The population of Ara-3 at generation 33,000 should be similar enough to the population of Ara-3 at the experiment's inception for us to judge if the evolution connecting them can be characterised as progressive. The answer here might seem to be 'yes': the later population of Ara-3 is able to do something which the earlier population of Ara-3 was not, it is able to metabolise citrate and so is 'better' at this ability than its predecessors. But this would be wrong, at least in the context in which we supported the use of term 'better' to describe functional improvement as in chapter three.

For the Ara-3 lineage to show improvement in this respect the ancestral population would need to metabolise citrate and the later population would need to metabolise citrate to higher- or 'better' - standard. This has not happened here - or at least has not happened yet. As it stands, the comparison between the non-citrate metabolising ancestral population and the citrate metabolising descendant population more closely resembles a comparison between the whale's flippers and Pakicetus' legs more than it does a comparison between the speeds of Olympic sprinters. The Ara-3 lineage features populations who differ in their *possession* of a given ability, not populations who feature differences in their *performances* of a given ability; to be progressive the evolutionary process must make later populations

perform the same selected effect to an improved degree. But notice that as events in evolution can be reappraised through the perspective of their future effects, the ‘progressiveness’ of citrate metabolism in the Ara-3 lineage would become clearer should its evolution *past* 33,000 generations feature improvements in citrate metabolism. Then we would *reinterpret* the advent of citrate metabolism thus far in the lineage previous to generation 33,000 as the first stage in a progressive sequence of functional improvement, and according to the selective effects conception, would be quite justified in doing so. This is why our current attitude to the ‘progressiveness’ of the Ara-3’s evolution to date must remain ambivalent; we do not yet have *enough information* about the effects of the potentiating and switching mutations for us to view them as progressive, but neither do we have the required information for us to rule their progressiveness out: Caesar has crossed the rubicon- so to speak - and we are waiting to see what happens next. If the Ara-3 lineage were to further evolve improvements in its citrate metabolism after generation 33,000, then the evolution preceding generation 33,000 will be the start of a progressive sequence; if the lineage does not then improve its citrate metabolism then while the mutations are important we would not then think of them as progressive in the same way; despite the mutations themselves remaining the same in either scenario.

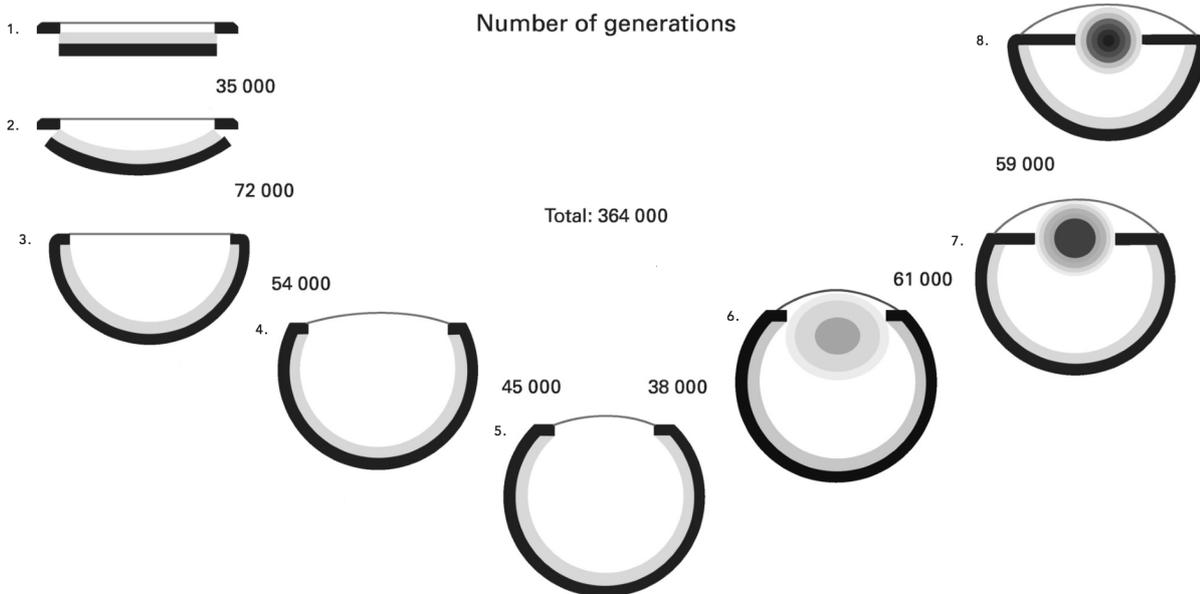
5.5 Nilsson & Pelger and the evolution of the eye

Designed to illustrate the relative speed and ease with which natural selection could develop a ‘fully formed’ eye despite what the authors describe as ‘pessimistic’ parameters, Nilsson & Pelger’s eight-stage linear depiction of eye evolution (Nilsson & Pelger 1994) has clearly been designed in accordance with the principles we identified above.

As with the the Ara-3 lineage, the sequence presented by Nilsson & Pelger arrives pre-delineated. We do not have to select a contemporary eye out from those possessed by a wide diversity of sighted species and follow the line of inheritance back to its distant ancestors, the lineage arrives to us with this procedure having been carried out. Nilsson & Pelger point out that the later stages of the eye are intended to be ‘vertebrate generic’ rather than an

anatomically accurate picture of an eye belonging to any vertebrate species in particular, so the picking out and tracing back of the lineage is a metaphorical rather than actual procedure as in the Ara-3 lineage.

Fig 5.2³⁸



The lineage presented is also quite clearly pre-demarkated, although in contrast to the Ara-3 lineage the number of stages dividing the lineage is very limited indeed: x8 in comparison to Ara-3's original x66, although the duration over which the evolution has occurred and changes made to the lineage depicted are both significantly greater. The motivation for the limit in both cases however is practical. For the real-world experiment of the LTEE, we assumed that the samples demarcating the lineage were taken at every x500 generations as opposed to a lower and more frequent amount due limitations in time, space and cost; for Nilsson and Pelger it is the requirement of presenting their findings in a manner that abstracts from the technicalities of their calculations and is straightforward to understand. With this objective achieved, the lineage does not need to be edited down further.

³⁸ Fig 5.2 (A simplified version of the original illustration in Nilsson & Pelger 1994 taken from Land & Nilsson 2002 and further reformatted - by Tim Ashton - for the purposes of this dissertation.) The lineage is demarcated into eight stages from 1 (top left) through to 8 (top right) with each representing a distinct stage in the eye's evolution. The figures between each stage are the number of generations the authors estimate each stage would take to evolve from that which precedes it, given what the authors believe to be realistic parameters

The first stage of Nilsson & Pelger's depiction is a flat patch of light-sensitive cells, backed by a darker layer and covered by a transparent cellular layer. So even at this early stage we have three types of cell and their specific configuration. Notice that starting the project here was a decision of the authors and was not naturally provided by the lineage itself. The project could have commenced earlier, the organ at stage 1 clearly did not arrive in the lineage fully formed, so those sceptical that a fully-formed eye could evolve over a mere 364,000 generations might be justified in asking for the lineage to be extended back further into evolutionary history so as to calculate the time taken for the development of each of these cell types and how they could have become layered in the manner described. Given that Nilsson & Pelger's stated aim is to estimate the length of time a complex eye takes to evolve, starting from what is clearly a rudimentary 'eye' of sorts may strike us as being something of a short-cut.

The final stage is a generic eye similar to those found in contemporary vertebrates. Notice again that this final stage, like the depiction's first stage, is not provided naturally by evolutionary biology either but is again a conscious decision. The authors could have chosen to have halted the lineage sooner at the slightly less 'refined' stage 7, or in fact continued beyond that of stage 8 if Nilsson and Pelger were to consider the 'degenerate' organs found in formerly-sighted vertebrates living in lightless environments where vulnerable eyes are at a selective disadvantage. Degenerate eyes are of course just as much a product of evolution as are the organs depicted in stages 1 through to 8, although given their aim of estimating the time taken for 'eye evolution' if the authors had extended their depiction beyond this most anatomically-complex organ it would both increase the evolutionary time required and all for a less-compelling result.

Notice also that as estimating the length of time or number of generations required for evolution to 'bridge the gap' spanning stages 1 and 8 is the aim of the project, these measures could not be the means by which the intervening stages were selected³⁹. As

³⁹ Nilsson & Pelger's explicit aim is to estimate the number of generations, rather than length of time, that a vertebrate eye could be produced by the evolutionary process. But as they stipulate generation time to be an average of one year over the length of the lineage in question, time in years or number of generations can be used interchangeably.

Nilsson and Pelger do not know when the ancestor of the vertebrate at stage 8 possessed the rudimentary 'eye' resembling that at stage 1 they cannot use that information to straightforwardly divide the lineage up by the number of stages required. Instead, their approach is to demarcate the eye's transformation using what we can recognise as key stages in its incremental evolution from a fairly rudimentary layer of cells to the sophisticated organ at the lineage's close. Notice again that although the development of the eye appears to be more or less monotonic, this is an artefact produced by the differential amount of generations used to separate out each of the different stages. The evolution from the flat triplicate cell-layer 'towards' the fully formed eye at stage 8 only appears to be even and consistent due to the generational gaps being increased or reduced so as to produce this effect. The longest generational gap (at 72,000) is in fact over twice as long as the shortest (at 35,000), and the specific demarcation of the sequence into the stages shown produces a standard deviation of 12,351. If Nilsson and Pelger's calculations for the evolution of the vertebrate eye are broadly correct the process was considerably more uneven and erratic than their illustration suggests.

In common with Edit 2 of our whale lineage and the Ara-3 lineage, Nilsson & Pelger's eye lineage is edited with the express purpose of explaining how the stages later in the delineated sequence evolved from those preceding them. The whale's fully aquatic physiology, the Cit+ phenotype, and the 'fully formed' eye are where the investigations into the evolutionary process responsible for these adaptations begin. In common with these examples, the procedure of editing the larger amount of generations spanning the length of the lineage down into those presented was done with this explanatory goal in mind.

The dark, light-sensitive and transparent cell layers and their relative positions to each other at the 'front' of the eye remain constant but major changes are shown to occur between each of the eight stages, these include: deeper invagination and thus increase in surface area of the light-sensitive layer; the partial constriction of the aperture, which decreases the angle of light hitting the light-sensitive cells and improves resolution; and the later development of a lens to allow the eye the ability to focus. While the morphological changes displayed in each stage may initially strike us as quite a disparate group lacking a consistent property

that runs through the sequence from start to finish, like the disparate anatomical changes demarcating the whale lineage the overarching theme of the sequence can in fact be construed quite narrowly. For the whale lineage the various anatomical changes could be jointly explained by the lineage's gradual adaptation to a fully-aquatic environment, for Nilsson and Pelger's sequence it is the requirement that later stages successively-produce a contemporary vertebrate eye. The authors explicitly state that a guiding principle for the sequence depicted was its evolutionary plausibility; each stage should provide a selective advantage to those preceding it while remaining consistent with evidence from comparative anatomy. However, in contrast to the demarcation of the whale lineages above, each stage must increase fitness by improving a *single function only* (ibid. p53) which Nilsson and Pelger characterise as 'spatial resolution' or 'visual acuity'. Later stages must not only be selectively advantageous to their predecessors, they must be selectively advantageous *because* they improve the organ's ability to perform this narrowly-defined function.

Once we appreciate the lineage's evolution is intended to be driven according to incremental improvements in a single function many of Nilsson & Pelger's choices become clearer. The authors begin their evolutionary sequence with the 'short-cut' of a triplicate layer of light sensitive cells as there is little guarantee that any of the organs ancestral to this stage could be characterised as driven by selection for 'visual acuity'. It is highly likely that before the lineage arrived at this fairly well-developed stage the photoreceptive cells evolved by virtue of alternative abilities; the synchronisation of internal processes to coincide with the day/night cycle being one plausible suggestion (Gregory 2008). Starting their depiction earlier in the lineage so as to include the possible evolution of these cells and their configuration would result in the sequence no longer holding together thematically; improvements in the function of visual acuity would no longer cover the reasons for each stage's advantage over its predecessors. Notice that this further explains why extending the lineage *past* the canonical stages of 1 to 8 to the 'degenerate' eyes of now-sightless vertebrates would also be inappropriate. The evolution of these auxiliary stages would not be driven by the fitness advantages of their increased visual acuity either when compared to the fitness of the stages preceding them because the visual acuity of organs at these later additional stages would be considerably worse.

So the delineation and demarcation of the sequence is provided, and Nilsson and Pelger's reasons for presenting the evolution of the eye in this way unambiguous. There is no requirement to further edit the lineage to a smaller number of stages as these are already quite brief. Again the stages 1 through to 8 certainly *appear* progressive; at the start of the depicted lineage we have some fairly 'rudimentary' layers of cells which over time, and thanks to the effects of natural selection, incrementally evolve into the complex and intricate organ depicted at stage 8. But unlike the radically-transformative yet non-progressive evolution of the whale-lineage, and unlike the ambiguous but potentially-progressive nature of the evolution of citrate metabolism in the LTEE, the evolution depicted here is unambiguously progressive in fact. As the consecutive stages of the eye's gradual evolution were explicitly chosen because they chart successive improvements in a narrowly-defined function, the organs later in the sequence will be better at that function than those preceding them. Later stages do not merely depict 'better' eyes, in the way that the later stages of the whale lineage could be characterised as depicting 'better' whales, the organs at the later stages in Nilsson and Pelger's lineage are functionally superior to those which precede them. So although the lineage *looks* progressive and the lineage *is* progressive this is due to *different* reasons. The anatomical changes depicted in Nilsson and Pelger's lineage look progressive because they incrementally move the lineage closer to the anatomy of the contemporary eye, but these changes are progressive because they successively produce increases in visual function in comparison to the organs of their predecessors. It so happens that in Nilsson and Pelger's lineage, rather than in that of the whale or in Ara-3 that the two qualifications coincide. The artificial spacing of the stages may lead us to think that the progression 'towards' the eye was smoother and more even than the sequence depicts, but as long as the function of visual acuity of each stage is increasing, the progressive nature of the eye's evolution remains intact.

5.6 Cross-species functional improvements

As we saw in chapter 3, should we identify the function of a trait as the effect or effects for which it provided a historical selective advantage as per the SE conception, then what the function of a trait is will be a matter of fact, and traits with the same or very similar histories of selection will have the same function. This provided sufficient grounds - the constitutive norms - for a principled comparison of their functional performances to be carried out. But our warrant for comparison only stretches so far; the selected effects conception of biological function provides an effective means for the comparison of *traits* only, and only then when the effects for which the traits were selected are comparable. It does not allow us to compare whole organisms with other organisms or compare species with other species, this is no sacrifice on our part, but a limitation which accords with our - and Darwin's - suspicions that attempting to compare entire organisms of very different kinds would inevitably produce results which were hopelessly impressionistic.

Although we have no grounds for a straightforward comparison of 'species with species', the SE approach to function will in certain circumstances allow for a comparison of the functional performances of the *traits* of one species with the functional performances of the *traits* of another species. This will be possible if the species in question are closely related and the traits in question inherited from a common ancestor, as up until their divergence the selectively beneficial effects of the traits would be identical. This can again be shown more clearly using a solid example.

Despite being the traits of different species, we might think of the claws and teeth of red squirrels and those of grey squirrels as performing the same functions in both species given their comparatively recent phylogenic divergence and the similarity of their environments since their separation occurred. Clearly in circumstances like these where the traits in question are closely related, the effects for which the traits were selected have a much greater chance of being the same or very similar effects and so permit their comparison; asking which has the 'better' teeth, a red or a grey squirrel would be a reasonable question if the tasks demanded of the teeth of red and grey squirrel teeth were the same tasks, and one

which could be answered if one species on average performed those tasks to higher standard. Conversely, the more distantly-related any two species are, the *less* warranted a comparison of their traits appears to become; at least it could seem this way if we were to think that the SE conception were the *exclusive* means by which such abilities could be identified. The traits of two grey squirrels - members of the same species - will have identical functions as the historical effects for which those traits were selected are identical; the effects for which the traits of more distantly-related species such as those of the red squirrel and those of the grey squirrel were selected may be very similar - and perhaps identical - but could be different to some degree due to differences in their selected effects; while the selected effects of traits belonging to very-distantly related species such as the grey squirrel and the North American beaver would seem to be too distinct for any comparison based on their functional performances to be carried out. The effects for which beaver teeth provided a selection advantage and the effects for which grey squirrel teeth provided a selection advantage will quite clearly be different effects. North American beavers will have better or worse North American Beaver teeth, and grey squirrels will have better or worse grey squirrel teeth (or perhaps simply 'squirrel teeth' - should red squirrel teeth share the grey squirrel teeth's functions). But beavers do not have 'better' teeth than squirrels, they do not perform the functions of squirrel teeth to a higher standard, and squirrels do not have 'better' teeth than do beavers for the same reason. When the traits of different species have different functions, the terms 'better' or 'worse' will not apply.

Notice however that in claiming that natural selection is likely to divergently-adapt the traits of formerly-unified species past the point where a comparison of their functions can no longer be carried out, we seem to present one of this chapter's central findings with a significant problem. While the SE approach provides *sufficient* theoretical grounds to warrant a comparison of the functional performances of traits belonging to the same or closely-related species, belonging to the same or closely-related species cannot be a *necessary* qualification for a valid comparison of the functions of traits to be carried out. This latter claim concerning the exclusivity of the SE approach is one that we must reject, given that our paradigm example of functional improvement - the increases in visual acuity as charted by Nilsson and Pelger's eye lineage - is concerned with improvements in a *single*

function evidently spanning the evolution of traits possessed by *multiple* species, albeit those belonging to single lineage. Whichever species possessed the triplicate layer of cells at the start of the eye-lineage is clearly distinct and very distant to the contemporary vertebrate which possesses the recognisable eye at the lineage's close, and yet we maintained that a functional comparison of the organs possessed by the various species *throughout* this lineage can be objectively carried out.

So on the one hand we are stating that the cross-species lung-function of Pakicetus and those of the modern whale are too different to warrant a comparison in terms of 'better' or 'worse' but consider the cross-species eye-function of the organs in Nilsson & Pelger's lineage to be warranted, although the anatomical differences between the earlier and later lungs of the Pakicetus-to-whale lineage are arguably less than the earlier and later organs of the eye lineage. How do we justify this apparent double standard?

This problem rests on a subtle conflation, while the SE conception defines what: i) the *function of a trait* is, it does not define what: ii) the *abilities which are functional* are. As we saw in chapter 3, the SE approach tells us that out of all the effects which an organism's given trait has produced, it is those effects which historically promoted the survival and reproduction of its ancestors which are that trait's function or functions. For the eye lineage this effect is visual acuity, for the lungs of the Pakicetus it is the exchange of certain gasses given certain environmental conditions. As we can see, whatever i) that function is - whatever those selected effects are- it is not *merely* an effect with a reliable history of the promoting the fitness of the organisms possessing it, it is additionally ii) a characteristic type of interaction between the organism and its environment - which then provides a selection advantage. The property of i) and the property of ii) are conceptually distinct and can come apart. As we appreciated, two traits with the same selection history will by necessity have the same functional abilities as per the the SE conception, but to have the same functional abilities two traits do not *need* to have the same selection history, it is not necessary for the selection histories of traits to be the same or similar if there are independent means of assessing that the effects which those traits produce are the same

effects. In the case of the eye lineage these means are available, in the case of the whale pakicetus-to-whale lineage they are not.

Let us unpack this claim. Consider the patterns on the wings of the Viceroy butterfly: these mimic the patterns on the wings of the Monarch butterfly, and as predators tend to avoid Monarchs, due to the similarity in their wings' appearance these predators will tend to avoid Viceroy's also. This is the effect for which the patterns on the wings of Viceroy's have been selected and so according to SE conception are the wing-pattern's function. The Viceroy wings' history of selectively advantageous effects tells us what i) *the function of their pattern* is: it is mimicry, but it does not define what ii) *the ability of mimicry* itself is: mimicry is the ability of one species to appear as though it were another or as though it were an inert feature of its environment. In the case of the Viceroy, the wings pattern's: i) *selected effects*, and: ii) *ability to resemble a member of another species* coincide, but they are clearly not equivalent properties and are conceptually distinguishable. The selected effects of the wings *could* have had nothing to do with mimicry - their patterns attracting mates instead to give a plausible example - and the Viceroy's ability to resemble another species *could* have proven to be a selectively *disadvantageous* effect- if it hindered the butterfly from attracting potential mating partners.

So when we agree with Nilsson & Pelger's assessment that 'visual acuity' is the function which increases over the eye lineage's depicted span and which provides each successive stage with a selective benefit over those which preceded, we are in fact supporting *three* separate claims. The first claim is that a) *each stage in the eye lineage is selectively beneficial to those which preceded it*, the second claim is that b) *the eye's ability to detect light and produce a sharper image- its visual acuity - increases as the lineage progresses*, and the third claim is that c) *the effect which makes each stage selectively beneficial to those which precede it, is the eye's increasing ability to detect light and produce a sharper image*. What the function of the eye is - *the effects for which it was selected* - and what visual acuity is - *the ability to detect light and to produce an image* - coincide in the eye lineage, but the relationship here is contingent. The *ability to detect light and to use it to produce a visual image* is a selected effect, but it is not *merely* a selected effect: it is a complex ability

produced through a range of physical and anatomical factors: the physical properties of light itself, the area of light-sensitivity, the refractive index of the medium through which light travels and so on. All of these conspire in the eye lineage to create an effect for which the eye is subsequently selected, visual acuity improves over the lineage because greater visual acuity is selectively beneficial, but being selectively beneficial does not tell us what visual acuity consists in. As long as we can *identify* the various abilities produced by traits - be it mimicry or visual acuity - independently from their histories as selected effects, there should be no prohibition on the comparison of the performances of the traits possessed by different species, as long as the abilities produced by those traits are potentially measurable like mimicry, and indeed, visual acuity.

It is not differences in historical selection pressures *per se* which prevent cross-species functional comparisons being carried out, but the fact that differences in selection pressure will very often result in the abilities of traits diverging past the point where the ancestral and descendant traits will be performing the same functional ability. This does not happen in the eye lineage as we can recognise the same ability of visual acuity in the triplicate layer of cells at the start of the lineage through to the fully-formed eye at the lineage's close- thus permitting their comparison. This consistency in ability almost certainly does not occur to the traits over the Pakicetus-to-whale lineage and thus prevents their comparison. Kidneys of course remove waste from the blood stream and lungs extract oxygen from the air and expel carbon dioxide whether they are in the bodies of pakicetus or in the bodies of whales, but what counts as 'good' or 'efficient' waste removal or gas exchange is highly dependant on the rest of the organism's physiology in the way the visual acuity is not.

So while the SE conception warrants the comparison of the functions of traits belonging to the same or closely-related species; the functional performances of the traits of species which are not closely-related can still be carried out if the ability is not 'species specific' in the way the lung and kidney functions of pakicetus are, but the increasing visual acuity over the eye lineage is not. The SE conception provides assurances that when we do compare the functional performances of closely-related traits that we are comparing the natural properties of those traits - the effects for which they were selected - and not properties that

are imposed on them artificially by observers like ourselves. As we saw in chapter 3, the SE conception provides the constitutive-norms permitting appropriate comparison. But whatever the functions of traits are, they are not *merely* effects which have been selected: the function of the eye (at least those in the lineage of Nilsson and Pelger) is to detect and produce visual images; this is the function of the triplicate layer of cells at the start of the eye-lineage, and is also the function of the fully-formed eye at the lineage's end, despite these organs belonging to very different - albeit distantly related - species. As each stage in the lineage is successively better able to perform this function - which is something that can be independently measured - and as each stage evolved because it was better able to perform this function, the evolution depicted in the eye lineage is progressive. Later eyes are *improvements* on those which precede them.

5.7 Chapter recap & conclusions

In section 5.2 we identified *delineation* as a valid operation: lineages are objective parts of the natural world, and treating them as such should not be controversial. We do however need to provide good reasons to prioritise certain lineages over others particularly when the lineage in question is that of humans beings - although this does not rule the prioritisation of the human lineage out if alternative justifications can be given. In 5.3 we discussed *demarcation*. As lineages can be produced out of countless individuals and many species, the objective nature of the 'stages' through which a lineage evolves are less straightforward. Lineages can be demarcated in many different ways and often the reasons for dividing a lineage will be so as to characterise the lineage's evolution as a 'progressive' process; changing the characteristics of the organisms and species earlier on in the lineage so that they resemble the characteristics of the organisms and species at the lineage's close. Once the 'theme' bridging an ancestral and descendant population has been decided on there will be broad agreement on which stages of the lineage should be emphasised and used to demarcate any evolution that has taken place and which should not. In 5.4 we saw that although evolution is not forward-looking, the procedure of demarcation is. When the demarcation of a lineage is done so as to explain the incremental evolution of the lineage's

later stages, the appearance of progressive evolution towards the lineages later stages will be the result. As we realised in 5.5 however, sometimes reality and representation coincide; a lineage can look progressive for the reasons stated in section 5.3, but actually be progressive for different reasons as stated in chapter 3; later stages of the lineage may be functionally superior to those which proceed them. In 5.6 we were initially concerned that our equating of evolutionary progress with functional improvement and our reliance on the SE conception of function revealed a potential weakness in our approach: it would seem to prevent the comparisons of the functional performances of distantly-related species, and so severely limit what counts as evolutionary progress as a result. We observed however, that although the SE conception is sufficient to tell us what the function of a trait is, it does not define what the abilities which are functional are; it is possible to identify these in different species independently of the SE approach.

We conclude this dissertation in the chapter ahead by reviewing and reappraising the findings of previous chapters. In our introductory chapter we posed some *prima facie* possible objections to the project of defending progress in evolution, highlighted four underlying themes whose influence would guide our analysis, and acknowledged three ‘central properties’ endemic in the work of previous authors’ attempts to defend progress in evolution with some degree of formality. We revisit these points with the benefit of our newly informed perspective.

Chapter 6

Evolution, progress and biological function: review and conclusions

6.1 Three objections reappraised

The valedictory tasks of this chapter are to recap, reappraise and conclude. In our introductory chapter we signposted four prominent themes in the work ahead: the explanatory scope of natural selection, the applicability of normative terms to natural phenomena, the relationship between representation and reality, and the methodology used in understanding evolutionary phenomena. We also identified three ‘central properties’ that previous accounts of evolutionary progress have implicitly respected: direction, drive and improvement. Additionally, we were alerted to distinct objections from three separate camps: empirical science, naturalist philosophy, and those concerned with the adverse ‘political’ implications a defence of progress in evolution might unwittingly support. With the benefit of the preceding analysis we are now in a position to answer these objections, discuss these prominent themes, and recognise how well our perspective on evolutionary progress respects the ‘central properties’ of those preceding - as well as extending and improving on them. We will start with the objections, concurrently recapping on this dissertation’s main points as we do so.

Our investigation began on the ‘back foot’; defending the casual efficacy and creative abilities of natural selection against a range of attacks designed to characterise it as a statistical abstraction or wholly eliminative influence. We remained unconvinced by these arguments and noticed that despite their differences they were based on a shared misconception: that events happening to individual organisms - whether by their proximate environmental interactions or by the specifics of their genetic inheritance - are the exclusive sites of evolutionary causality. We emphasised that natural selection’s causality cannot be directly witnessed at the individual-level, but that this did not make natural selection at the population-level derivative. Individuals are parts of populations, therefore an individual’s

environmental interactions *are* a population's environmental interactions, not derived from them. We then emphasised that although the creation of novel traits are due to individuals' genetic inheritance and subsequent development, it is a mistake to malign natural selection's influence in this process; it plays a crucial role in explaining why individuals have the specific genetic inheritances that they do. Were it not for natural selection retaining unlikely but selectively advantageous variation in a individual's ancestral population, complex traits which require the combination of many such variants would be vanishingly unlikely to ever evolve.

As we were defending the powers of natural selection that orthodoxy already recognises it to possess, our thinking on this subject spoke to none of the three camps' objections directly but rather provided a secure footing for our later progressive claims.

We then reappraised the philosophical problem of providing naturalistic warrant for the use of normative terms. The SE conception of function proved to be an essential tool in this project; through it we were able to naturalise biological function, and through naturalising biological function we were able to provide sufficient, empirical grounds for describing the functioning of specified traits as 'better' or 'worse' than the functioning of others. We very much moved from the back foot into a proactive stance here: our progressive interpretation of natural selection operating through differential functional effectiveness elucidated the sometimes confusing, yet widespread, practice of explaining the existence of complex traits due to them being 'selected'. We observed that this manner of description is frequently elliptical: what is fully meant is that the causal history of complex traits is often one of trait variants successively providing a selection advantage due to increases in their functional efficacy: of traits *improving*.

We can see clearly here that the scientific camp's objection that natural selection should not be thought to produce improvement in any way is explicitly addressed. Our use of improvement is not supplementary to empirical explanation or conjured by conceptual sleight-of-hand, but an inherent part of the explanation for how complex adaptive traits are produced. Saying that trait A has a selection advantage over trait B may be an *adequate*

explanation for trait A's prevalence over trait B, but it is far from a *comprehensive* explanation: for that we need the empirical facts on the ground, and that will involve knowing what trait A and trait B are doing. As we noted, perhaps whatever advantage A provides is very different to that of B, but our numerous examples where selection advantage was provided by functional improvement were not hard to come by; we can easily recognise this dynamic to be quite general and widespread.

Notice that warranting the use of 'improvement' empirically should also put the philosophical camp's concerns to rest as well as those of empirical science - at least in part. If the concepts of 'improvement' and 'betterment' perform an important explanatory role in evolutionary science as we have claimed, it is surely not for naturalist philosophy to intervene and counter-claim that this mode of explanation is invalid and its results therefore wrong. As we have stressed, the application of normativity here is: i) of a very different stripe to that concerning the actions and intentions of moral agents and about which philosophy quite rightly has much to say, and: ii) accurately describes genuine causal interactions. Our darker butterflies in their industrial environment proliferated because they were better camouflaged, later eyes were selected because of improvements in their visual acuity; if naturalist philosophy feels that this is an inappropriate characterisation, we might at this point suggest that the onus is on those opposed to explain why.

This answers the philosophical naturalist camp's normative concerns, how does our perspective fare in the light of potential anti-naturalist criticism? Notice we seem to be very far away from the top of a 'slippery slope' into anti-naturalist rhetoric here. Having grounds for the use of 'improvement' gives us grounds for the use of 'progress'; a lineage featuring traits selected by virtue of their functional improvement will be a 'progressive' one in respect of that specific improvement⁴⁰, and we could, if we wished, describe better functioning traits as more 'advanced' or 'superior', perhaps even 'better designed' -

⁴⁰ Although a lineage could of course *not* be progressive in respect of other traits which are not improving in their functional efficacy over the same duration - or in respect of traits which are getting functionally worse.

although we will return to the issue of ‘design’⁴¹ more fully below. What ‘improvement’ does not provide is warrant for thinking that nature is teleological, purposeful, or manifests intent in any way. Again we must credit the SE conception of function for granting us immunity to this criticism. As is well established, a fundamental strength of the SE conception is its ability to account for the *apparent* purposefulness of traits naturalistically. Functions are not the goals of traits, they are not what traits are *for*, functions are past effects for which traits were selected, the mode of explanation make no reference to goals, purposes or objectives and our progressive use of the SE concept of function does not introduce them nor provide any grounds for this.

Not wishing to consider our confirmation of natural selection’s progressive credentials to be the exclusive means by which progress in evolution could be justifiably defended, we outlined what qualities ‘big picture’ evidence of progress - that showing improvement across species, in higher taxa, or perhaps even in the history of life as whole - would need in order to be considered authentic and representative of genuinely progressive phenomena. We identified the essential ingredient for ‘big picture’ progress as a mechanism for the phenomena, and further suggested the properties we would expect a mechanism for the generation of authentic large scale evolutionary phenomena to possess. Lacking these properties Wilson’s pattern of ‘diversity-promoting’ stages was found to be quasi, whereas Dawkins’ pattern of co-evolutionary arms races was found to be authentic - and moreover somewhat familiar. The incremental improvements in function we identified as the basis for evolutionary progress are likely to be a major feature of the arms races Dawkins describes.

After widening our perspective to cross-species and higher taxa we then narrowed our focus in order to examine the evolution which occurs over evolutionary lineages: the procedure we use to delineate a single lineage out from the ‘phylogenetic bush’ and the procedure we use to demarcate the characteristic stages of the evolution which the lineage has undergone. We noted that the recognition of an evolutionary lineage is very different from the

⁴¹ We will be keeping the term ‘design’ in inverted commas for the remainder of this chapter. Dennett prefers to capitalise Design (Dennett 1995a), while Dawkins at times uses the term ‘Designoid’ (Dawkins 1997) - both perhaps wary of deliberate misinterpretation. To my thinking little hangs on this.

prioritisation of an evolutionary lineage: something which has been conflated in the claims of anthropocentrism often held by those sceptical that evolutionary progress amounts to anything over and above human aggrandisement. While the recognition of an evolutionary lineage is uncontentious, the manner in which its evolution is represented is open to considerable interpretation. We again considered some examples: the evolution of the whale while clearly radical, was not progressive; whereas the evolution of the vertebrate eye was progressive. Once more the reasons for this were familiar: the eye lineage charted a clear series of improvements in a narrowly defined function in the manner spelt out in chapter 3, and in the manner that the offensive and defensive traits of the species locked into Dawkins' arms races are likely to improve.

Here again we can observe two camps' concerns. Although this time it is the political camp's unease of differentially apprising the natural world, combined with the philosophical camp's lingering worries over the reintroduction of teleology.

As we noticed in our investigation of demarcation, these philosophical worries are not unfounded. Charting the steps towards the 'end state' of an evolutionary lineage - whether the Cit+ phenotype or the vertebrate eye - in order to understand how it was achieved is the express purpose of this type of evolutionary explanation and is inherently teleological. As we appreciated however, the teleology is a property of the method, not a property of the phenomena. We understood the evolution of the vertebrate eye as a progressive process - and not a teleological one - because over the span of the lineage the organs depicted were functional improvements of those which preceded them - not because the 'final' product was in anyway predestined. The procedure of demarcation and its usual application may be misleading but our approach to evolutionary progress is quite separate and distinct from this.

The concern that our limited defence of evolutionary progress may provide cover for 'human imperialism' can also be dismantled. Exposing the flaws in an argument which conflates progress with anthropocentrism does not in turn imply a defence of this position. As we learned, progress in a very real sense can be found in the improvement of traits

possessed by organisms far removed from humankind, from primates, from mammals, and indeed even from eukaryotes. Once more the SE conception provided a welcome break on unwarranted conceptual expansion: it is functions which are performed better or worse, and it is traits which have functions - then it is organisms which have traits. We can say that trait A is better than trait B *in the context of* performing function *f* by virtue of differences in A's and B's functionality, but the most this tells is that A and B as *traits* are better or worse, this does not give us grounds to think of the organisms possessing A and B as better or worse *organisms*; the evaluations of trait performance are not transferable to organisms in this way.

If we still find the question of specifically human evolutionary progress to be of interest we might consider the following. Firstly, it is unclear what functional improvement has occurred in any great degree over the human lineage, let alone is *characteristic* of it. What selected effects do we think later members of the human lineage perform to a higher standard than earlier members? We may, if we wish, provide our own suggestions to this question adding to what is already a contested and over-crowded field. However secondly, if human evolution was progressive, the most that this will tell us is that some of the traits of later members of the human lineage are functional improvements over the traits of earlier members; this does not remotely provide any platform for the notion of human superiority, nor does it provide even the thinnest veneer of scientific respectability for the abuses of other non-human members of the natural world.

6.2 Four central themes revisited

As we can now appreciate, the four central themes which we identified in our opening chapter were strongly represented throughout our analysis; we rarely strayed far from at least one of them.

Our main theme concerned the power of natural selection. This issue was addressed most explicitly in our defence of natural selection as both a causal and as a creative mechanism

but was pervasive throughout. Through doing this, we not only reaffirmed natural selection's causal and creative efficacy *per se* but were able to appreciate the *sort* of creativity that natural selection is able to produce: composite traits requiring multiple individual traits for their production. We additionally recognised that the history of a trait's selected effects informs us as to that trait's function - natural selection's role was somewhat different here in being diagnostic in addition to being casual. We then observed that natural selection was able to account for improvement: to the extent that the current effects of traits are of the same type as the effects for which they were historically selected, their current selection advantage will be due to their improved functioning.

We can notice the scale of the phenomena for which natural selection was held responsible increased greatly as our work progressed. We started with mere changes in population structure, then moved to the creation of multi-variant traits, then in our investigation into the authenticity of large scale evolutionary patterns we identified natural selection as the only suitable mechanism, although stopped short of suggesting that such authentic patterns exist. The explanatory scope of natural selection overlapped to a considerable extent with our second thematic thread: the application of normative terms to natural phenomena. Unsurprisingly, as our warrant for the use of normative terms was ultimately due to the powers of natural selection - albeit through the intermediary of biological function.

As we noted above: normativity's reach into evolutionary science was strictly limited. The possession of a function does not impose 'obligations' on traits as supporters of the SE conception of function have variously contended, nor does the SE conception provide unrestricted grounds for the application of normative terms outside of the narrow remit of biological function. If the function of a trait endures over extended spans of time and anatomical change then the normative appraisal of the functioning of that trait in terms of better or worse and so on will be appropriate. But this is only warranted due to the persistence of the function, if the nature of the effect for which a trait is selected changes - as we might expect it is very likely to do over extended spans of time - the comparing of later and earlier 'functioning' becomes incoherent; it would be unclear what the later versions of the traits were better *at*.

We saw the limitation on normative reach provided by the SE conception quite clearly exercised in Wilson's 'progressive' approach to biodiversity. Even if an underlying mechanism for Wilson's pattern - or indeed any pattern discerned in the evolution of life at scales above the functions of traits - were discovered, this would make the pattern authentic, but it would not make the pattern one of improvement. We may find biodiversity, biological complexity or hierarchical organisation to be of particular interest, and have very good reasons for valuing it, but if so we can reasonably ask for the basis for these normative attitudes. What are these values based on? The SE conception of function provides a clear answer to this, increasing biodiversity and complexity in themselves do not.

Notice that a consideration of our third theme: the interplay between representation and reality provides somewhat mixed results. The traditional and explicitly progressive Malthusian-based interpretation of species competing in the struggle for survival with natural selection driving the survivors to ever greater strength and sophistication proves to be approximately correct; the caveat being that natural selection does not *necessitate* progressive improvement as a naive interpretation of the 'struggle for survival' might assume. Organisms 'outcompeting' their rivals is not enough, competition must be in the strict sense that participants are competing at the same activity and superior functional performance in that activity won out. In contrast however other highly prevalent ways of representing evolution were shown to be potentially misleading. We have already revisited the shortcomings of the 'big picture' approach to large scale evolutionary phenomena, but in addition to this, the procure of demarcating an evolutionary lineage into its component 'evolutionary stages' invariably presents the nature of the actual evolution from which it has been selected as progressive, whether that evolution actually is progressive or is not. We rightly object to the 'onwards and upwards' depiction of human evolution as a series of stages leading progressively from quadruped ape to spear-carrying man, but as our Pakicetus-to-whale lineage demonstrated, other depictions based upon explaining later stages through their evolutionary predecessors provide the same impressionistically-progressive result; and as our analysis of the eye lineage's representation demonstrated,

even when the evolution depicted *is* progressive, this approach presents it as *more* progressive than it actually is.

Our fourth underlying theme concerned the methodology used in understanding evolutionary phenomena. This theme touched upon many of the others. Based upon methodological approach we might characterise our preceding work as composed from two halves. The first half - defending and reappraising the explanatory reach of natural selection - utilised a 'bottom-up', process-centred approach, in contrast to the second half which employed an approach explicitly 'top-down' and 'pattern-centric'. With the benefit of hindsight we can now recognise these approaches as complimentary. On the one hand, the process of natural selection need not produce any progressive phenomena - it need not produce anything at all - making an exclusively bottom-up approach based solely on natural selection insufficient. While on the other we need to see what phenomena natural selection is capable of producing - to witness its effects - in order to confirm that its influence is one of progressive improvement. As we have addressed already, a top-down approach based solely on the apparent 'progressiveness' of evolutionary phenomena is insufficient. Progressive-*seeming* evolutionary phenomena may be generated arbitrarily, or through causes for which the term 'progressive' would be inappropriate.

We might have cause to think that despite our claims to the contrary the explanatory reach and casual efficacy of natural selection has been given too great a prominence in this dissertation, that natural selection has acted as something of *deus ex machina*, and that we are adaptationists way beyond sensible limit. This characterisation would, I think, be unfair. We do not intend to malign the critical role that mutation plays by creating genetic variation and therefore its contribution in the generation of functional improvement itself, nor do we contend that natural selection is the only conceivable means through which large-scale authentic evolutionary patterns may be generated; as our pains to set out the required properties for an alternative mechanism should contend. Clearly however, as our warrant for the normative application of 'improvement' is based upon the SE conception of function, and the SE conception of function is explicitly based upon natural selection, our claim that

‘improvement’ could be accounted for without natural selection might ring hollow. But consider the following:

Recall *fig 5.2* - Nilsson & Pelger’s depiction of the evolution the vertebrate eye. We might be confident that if anything *looks* like an example of progressive evolution it is this: a sequence charting both increases in anatomical complexity and sophistication of ‘design’. We established that this series was indeed progressive as the organs’ selected effect - visual acuity - could be seen to improve throughout. This would be impossible to defend if natural selection were not the cause, as there would be no selected effect. As we also saw however, *what the function of a trait is*, and *the abilities which are functional*, are conceptually distinct and the SE conception of function only has purchase on the first of these properties. The ‘eye’ at stage 1 of Nilsson & Pelger’s sequence still produces a visual image - *the ability which is functional* - as does the eye at stage 8 no matter how these organs were produced; it is just that without natural selection, visual acuity could not be these organ’s function.

Now let us completely remove natural selection from the equation and suppose somewhat improbably that 160 years of biological science is mistaken, that natural selection is not a genuine evolutionary mechanism, and that Lamarck was correct all along; evolution is driven by the inheritance of somatic differences acquired over the lifetime of an organism and passed onto their offspring. Clearly visual acuity could not be the eye’s SE function in these circumstances, but would we still consider the sequence to be progressive? Although removing the SE conception from consideration would also mean discarding the entire conceptual framework we have based upon it, we might suspect the answer to this question would very likely be yes: the sequence of eye evolution is still charting increases in visual acuity⁴². It is of course difficult to predict how we might appraise philosophically the results of a mechanism which does not exist - let alone one as outré as Lamarck’s - but the lengthening of the necks of the giraffes so frequently used to illustrate Lamarckian evolution

⁴² Although of course without the diagnostic ability of the SE conception to inform us as to the organ’s function we would have to produce alternative means for justifying that *the* effect of visual acuity should be given preference over the many other effects that the eye’s ‘Lamarckian evolution’ would increase.

were no less progressive because they were not produced by natural selection, and we might guess that neither would the eye's evolution seem any less progressive were it driven by the inheritance of acquired characteristics.

Natural selection is the only mechanism for evolutionary progress which we know of, but our work here does not provide grounds for thinking it the only *conceivable* mechanism.

6.3 Direction, drive, improvement... and design.

We noted that despite substantial differences in discipline and philosophical outlook, previous approaches to progress in evolution implicitly respected three central properties: direction, drive, and improvement. As we can now appreciate, the most contentious of these properties - improvement - is something that our perspective of progress as specifically functional improvement accounts for, and moreover accounts for naturalistically.

Consider some of our examples of progressive evolution: the increases in functional efficacy as evidenced by the visual acuity of the vertebrate eye, by the mimicry of the Viceroy butterfly, and by the motility of the bacterial flagellum. Clearly the evolution of all three of these traits is directional. Again, using Nilsson & Pelger's illustration eye evolution as a rough guide (*fig 5.2*), if we were to randomly arrange each of its eight stages we would be able place them in correct temporal order with a high degree of confidence; we could be reasonably certain that this would also be possible for both the changes in appearance of the Viceroy butterfly and the evolution of the bacterial flagellum also. Information concerning the traits' various functions would not be required to achieve this; with the incremental appearance of new parts and the incremental disappearance or alteration of others, morphology alone would very likely prove sufficient to do the job for us.

As we have explained extensively, it is equally clear that the evolution of each of these traits is driven and not accidental. In common with almost all post-synthesis perspectives on evolutionary progress, we account for the changes in the eye, and in the butterfly and

bacterial morphologies through the actions of natural selection. Subsequent stages in each of these lineages are retained because they provide a selection advantage over those which precede them.

Concerning the property of improvement, this again has been explained extensively. We might however, emphasise that the improvement that we have defended in this dissertation is not only compatible with mainstream evolutionary practice and philosophical naturalism, but, plausibly performs an explanatory role for our understanding of how complex, ‘design-like’ traits are produced. The progressiveness of these traits’ evolution is not something imposed upon them by teleologically-minded or apophenic observers, the incremental construction of these traits is characterised by improvements in their functionality and are quite real. Improvement in these traits’ functioning is not some sort of explanatory optional extra - it has genuine casual efficacy. Later variants in such lineages are not merely selectively beneficial, they are selectively beneficial because they are improvements on those which precede them: later eyes are advantageous *because* they have improved visual acuity, later Viceroy morphology is advantageous *because* it provides improved mimicry, later bacteria enjoy a selective advantage *because* of their improved motility. A comprehensive empirical explanation for each of these traits’ evolution would clearly include natural selection to account for the ‘driven’ nature of changes undergone, but would be severely incomplete if the explicitly progressive nature of the ‘facts on the ground’ for each subsequent stage being selectively beneficial were not included.

We will sign off this dissertation with a conjecture: how our discussion on function and improvement might bear on the question of ‘design’ in nature - an aspect of adaptive phenomena which has so far remained somewhat in the the periphery of our project, avoiding direct engagement.

We noted briefly at chapter 3’s close that some improvements in biological functioning will be more easily achieved than improvements in others. The increases in camouflage enjoyed by the moths in the industrialised north for example could initially be achieved relatively simply, perhaps by the introduction of an allele making the organisms possessing it darker

and more difficult to spot; in contrast, increases in the nocturnal bird's auditory sensitivity will require more particular and precise modification and - *ceteris paribus* - therefore be harder for the accumulation of random variation to achieve. We might pause here and reflect on how differently we would apply the notion of 'design' or 'design-like' in each of these cases.

To fully bring these intuitions to light, consider how the concept of 'design' might apply to a lineage undergoing progressive evolution by virtue of variation producing improvements in functional efficacy. For the moth population lightly coloured wings will be badly camouflaged against the newly-darkened trees of their polluted environment; after the introduction of the 'darkening allele', the organisms possessing it become more effectively camouflaged than the population average. Given enough consistency in selection advantage we would say that the darker wings' colouration has the function of camouflage: for this would be the wing colour's selected effect. Now consider the arrival of another allele, making the moths' wings slightly darker. From our perspective we would understand this to be an improvement in their functioning, and so to demonstrate evolutionary progress to some degree. Do we also think that the colour of the darkest wings so far in this example shows evidence of 'design'? This seems like quite an odd question to ask, and one for which there is no clear answer. In terms of environmental suitability the wings seem 'more designed' than those lighter, but the notion here seems misapplied. The alteration to the trait strikes us as too *rudimentary* to warrant the use of the term.

Now consider a third allele introduced into the moth population, this time its phenotypic effect produces an even darker wing, increasing camouflage still further. Again despite this being progressive - a function is improving after all - we might still harbour doubts as to whether the language of 'design' would be appropriate. But if not now, then when? One suggestion for the partial incongruity between 'improvement' - which is happening here - and 'design' - which might not be - could be this: in lieu of an operational definition for the use of the term, 'design' seems to require changes to *more than one aspect* of the trait. If this is right, then the use of 'design' for the moths' darkened wing-colour would be referring to a single aspect and therefore seem somewhat pedestrian. Consider how we might appraise the

arrival of a different *kind* of variation in the moth population, perhaps an allele which softens the edges of the wings, breaking up the moths' outline and providing even better camouflage over those moths which lack it. Although improved, the identity of the function remains the same: camouflage, but now there are multiple different aspects of the moth conspiring to produce that function, 'design-talk' suddenly seems more appropriate. To drive this point home, consider additional variation in another phenotypic dimension of variation, increasing further the moths' camouflage ability, but on this occasion making their wings more similar in shape to the leaves of the trees of their polluted environment. Now with pattern, outline *and* shape conspiring to increase camouflage ability, the notion that the moth's wings exhibit 'design' seems near-unavoidable. The greater the amount of parts sharing what we might term 'functional integration' in this way, it seems the more their apparent 'design' impresses itself upon us.

Notice that this 'design-like' functional integration is exactly what is happening in our paradigm progressive example: Nilsson and Pelger's progressive illustration of eye evolution (*fig 5.2*). The flat triplicate layer of light-sensitive cells gradually deepens, then further along in the lineage, the aperture constricts, a lens is introduced then further refined. The dimensions along which change occurs in this sequence are many, but they are tightly integrated; both anatomically, and *functionally* in they together improve efficacy in the organ's given selected effect. While complexity *per se* is rightly dismissed as indicative of evolutionary progress, the 'design-like' qualities of traits such as vertebrate eyes, camouflaged wings, and bacterial flagella confer on them complexity of a recognisably different stripe. The parts of these traits are not merely diverse and distinct, they are diverse and distinct and *integrated* so as to produce a single functional effect. Once more our identification and use of function in our analysis is baring conceptual fruit. We might further suggest that we not only *recognise* 'design' in evolution through the contribution of diverse parts to a common function, but moreover suggest that we also *explain* the evolution of the 'design-like' properties of complex traits through the fact that, in some cases, increases in functional effectiveness may only be achieved through a more complex 'design' of the trait responsible for producing them.

This claim may be speculative, but consider that initially the increases in the moths' camouflage are comparatively easy to achieve, as darker colouration produces better camouflage the wings need only change in colour to increase their functional effectiveness. But as the moths can only match themselves to the colour of their environment up to a given point, once this has been achieved, for camouflage to increase further it will have to be due to changes in a different anatomical dimension; as we suggested above, perhaps the contours of the wings could be softened, or the shapes of the wings themselves modified. As each of these modifications provides a selection advantage, we are able to explain them through the mechanism of natural selection - obviously; but as we can recognise through this dissertation's analysis this is not natural selection *simpliciter*; this sequence is also *progressive* in that each subsequent modification is an improvement in the effectiveness of a given function- here camouflage. This might be an adequate description if functional effectiveness were increased due to change in a limited number of ways, such as when the wings become darker only. But when 'one-dimensional' change of this kind has reached a certain limit, functionality can only be increased by changes in a different phenotypic 'dimension', and it is when this occurs that the design-like character of complex adaptive traits becomes increasingly apparent. This is an increase in complexity, but complexity of quite a distinctive kind: the number of different aspects of a given trait which contribute to the production of a narrowly-defined function.

As is well known some, but not all, evolution is due to natural selection; a central task in this dissertation has been to demonstrate that that some, but not all, natural selection is due to functional improvement and will be progressive through doing so - a task we might feel that has been successfully achieved. Here it would seem that if complex design like-properties are the means through which the effectiveness of a given function can be further improved, progressive evolution - characterised as improvements in functional effectiveness - could explain how the design-like properties of complex traits come about. Traits gain design-like properties, as when one dimension of improvement is unavailable or exhausted design-like properties are an alternative means through which a trait's functional efficacy can be further increased.

If our thinking here is correct then in common with many alternative approaches to evolutionary progress this would lead us to consider the design-like qualities of traits as ‘advanced’ and indicative of evolution’s creative and progressive potential. We could be understood as following in the grand tradition of J.Huxley, Simpson, Mayr and so on - but with an important caveat. In contrast to these other approaches which consider such traits as advanced on largely impressionistic grounds, or perhaps due to anatomical complexity *simpliciter*, from our perspective this claim is empirically justified. The design-like qualities of complex adaptive traits are ‘advanced’ as they provide explicit evidence that progressive evolution has taken place; such traits exist *because* of their functional superiority.

6.5 Conclusion

Writing in 1988 the philosopher and evolutionary biologist Fransisco Ayala (in Nitecki (eds.) 1988) argued that: ‘...the notion of progress is axiological, and therefore it cannot be a strictly scientific term: value judgments are not part and parcel of scientific discourse, which is characterised by empirically testable and objective descriptions’ (ibid p81), before adding: ‘...no attempt to define progress as a purely biological concept has succeeded.’ (ibid p84).

In the light of the preceding analysis, I hope that it will now be clear in what sense Ayala was mistaken. It is true that the concept of progress will involve the application of ‘values’; whether the progress is societal progress, technological progress or the progress which we intuitively feel that evolution by natural selection is capable of producing. But Ayala’s ready classification of the diverse species of ‘value’ used in our understanding of nature as belonging to the single axiological genus was an assumption - and a wrong one. We have not only demonstrated that ‘value judgments’ *are* part and parcel of scientific discourse - as it is hard to understand how we might explain differential functioning without them; and we have also shown that the specific values involved in some evolutionary descriptions are empirically testable and objective - and moreover have both explanatory and causal power.

With these conceptual restrictions found to be unwarranted, the bar on appreciating progress in evolution as a purely biological concept has been lifted. Despite the interwoven histories of 'progress' and 'evolution' stretching back in antiquity, contemporary attempts to untangle these two concepts have been unsuccessful, and it seems, premature. Progress in evolution is real and moreover is a common feature throughout the tree of life.

Page left intentionally blank

Bibliography

- Adami, C., Ofria, C., & Collier, T. C. (2000). Evolution of biological complexity. *Proceedings of the National Academy of Sciences*, 97(9), 4463-4468.
- Archibald, J. D. (2014). *Aristotle's ladder, Darwin's tree: the evolution of visual metaphors for biological order*. Columbia University Press.
- Arendt, J., & Reznick, D. (2008). Convergence and parallelism reconsidered: what have we learned about the genetics of adaptation?. *Trends in ecology & evolution*, 23(1), 26-32.
- Ayala, F. J. (1970). Teleological explanations in evolutionary biology. *Philosophy of science*, 37(1), 1-15.
- Ayala, F. J. (1974). The concept of biological progress. In *Studies in the Philosophy of Biology* (pp. 339-355). Palgrave, London.
- Banks, W. E., d'Errico, F., Peterson, A. T., Kageyama, M., Sima, A., & Sánchez-Goñi, M. F. (2008). Neanderthal extinction by competitive exclusion. *PLoS One*, 3(12), e3972.
- Barros, D. B. (2008). Natural selection as a mechanism. *Philosophy of Science*, 75(3), 306-322.
- Beatty, J. (1984). Chance and natural selection. *Philosophy of science*, 51(2), 183-211.
- Beatty, J. (1995) "The evolutionary contingency thesis." *Concepts, theories, and rationality in the biological sciences* : 45-81.
- Beatty, J. (2006) "Replaying life's tape." *The Journal of philosophy* 103.7: 336-362.
- Beatty, J. (2008) "Chance Variation and Evolutionary Contingency: Darwin, Simpson, The Simpsons, and Gould." *The Oxford Handbook of Philosophy of Biology*.
- Beatty, J. (2016). What are narratives good for?. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 58, 33-40.
- Beatty, J., and Carrera I. (2011) "When what had to happen was not bound to happen: History, chance, narrative, evolution." *Journal of the Philosophy of History* 5.3 (2011): 471-495.
- Bender, R., Tobias, P. V., & Bender, N. (2012). The savannah hypotheses: origin, reception and impact on paleoanthropology. *History and philosophy of the life sciences*, 147-184.
- Blount, Z. D. (2016). A case study in evolutionary contingency. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 58, 82-92.
- Blount, Z. D., Borland, C. Z., & Lenski, R. E. (2008). Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. *Proceedings of the National Academy of Sciences*, 105(23), 7899-7906.
- Blount, Z. D., Lenski, R. E., & Losos, J. B. (2018). Contingency and determinism in evolution: replaying life's tape. *Science*, 362(6415)
- Bonner, J. T. (1988). *The evolution of complexity by means of natural selection*. Princeton University Press.
- Boorse, C. (1975). On the distinction between disease and illness. *Philosophy & public affairs*, 49-68.
- Boorse, C. (2011). Concepts of health and disease. In *Philosophy of medicine* (pp. 13-64). North-Holland.
- Bowler, P. J. (1978). Hugo De Vries and Thomas Hunt Morgan: the mutation theory and the spirit of Darwinism. *Annals of Science*, 35(1), 55-73.

- Bowler, P. J. (2009). Darwin's originality. *science*, 323(5911), 223-226.
- Boyd, R., (1988) "How to be a Moral Realist" in Sayre-McCord, G. (ed.), 1988. *Essays on Moral Realism, Ithaca and London: Cornell University Press.*
- Brandon, R. N. (1990). *Adaptation and environment* . Princeton University Press.
- Brandon, R. N. (1999). The units of selection revisited: the modules of selection. *Biology and Philosophy*, 14(2), 167-180.
- Brandon, R. N., & Rausher, M. D. (1996). Testing adaptationism: a comment on Orzack and Sober. *The American Naturalist*, 148(1), 189-201.
- Briggs, D. E., & Fortey, R. A. (2005). Wonderful strife: systematics, stem groups, and the phylogenetic signal of the Cambrian radiation. *Paleobiology*, 31(S2), 94-112.
- Brodie III, E. D., & Brodie Jr, E. D. (1999). Costs of exploiting poisonous prey: evolutionary trade-offs in a predator-prey arms race. *Evolution*, 53(2), 626-631.
- Brugger, P. (2001). From haunted brain to haunted science: A Cognitive Neuroscience View of Paranormal and Pseudoscientific Thought. In J. Houran & R. Lange (Eds.), *Hauntings and poltergeists: Multidisciplinary perspectives* (pp. 195-213). Jefferson, NC: McFarland.
- Brunnander, B. (2013). Did Darwin really answer Paley's question?. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(3), 309-311.
- Burnston, D. C. (2017). Real Patterns in Biological Explanation. *Philosophy of Science*, 84(5), 879-891.
- Butterfield, N. J. (2003). Exceptional fossil preservation and the Cambrian explosion. *Integrative and comparative biology*, 43(1), 166
- Calcott, B. (2008) "Lineage explanations: explaining how biological mechanisms change." *The British Journal for the Philosophy of Science* 60.1: 51-78.
- Calcott, B. & Sterelny, K. (2011) "The Major Transitions In Evolution Revisited" (Eds) MIT Press.
- Callaway, E. (2016). Oldest ancient-human DNA details dawn of Neanderthals. *Nature News*, 531(7594), 296.
- Carpenter, P. A., & Bishop, P. C. (2009). The seventh mass extinction: Human-caused events contribute to a fatal consequence. *Futures*, 41(10), 715-722.
- Cartwright, N. (2017). How Mechanisms Explain?. *Making a difference: Essays on the philosophy of causation*, 5.
- Catley, K. M., Novick, L. R., & Shade, C. K. (2010). Interpreting evolutionary diagrams: When topology and process conflict. *Journal of Research in Science Teaching*, 47(7), 861-882.
- Chalmers, D. J. (2012). *Constructing the world*. OUP Oxford.
- Christensen, W. (2012). Natural sources of normativity. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(1), 104-112.
- Clarke, E. (2010). The problem of biological individuality. *Biological Theory*, 5(4), 312-325.
- Cleland, C. E. (2011). Prediction and explanation in historical natural science. *The British Journal for the Philosophy of Science*, 62(3), 551-582.

- Conway Morris, S. (2000). The Cambrian “explosion”: slow-fuse or megatonnage?. *Proceedings of the National Academy of Sciences*, 97(9), 4426-4429.
- Conway Morris, S. (2003). *Life's solution: inevitable humans in a lonely universe*. Cambridge University Press.
- Conway Morris, S. (2006). Darwin's dilemma: the realities of the Cambrian ‘explosion’. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1470), 1069-1083.
- Conway Morris, S. (2008) “*The Deep Structure Of Biology*” (eds.) Templeton Foundation Press.
- Conway Morris, S. (2010) “*Evolution: Like Any Other Science It Is Predictable*”. *Phil. Trans. R. Soc.* 12;365(1537): 133-145
- Cooper, R. (2002). Disease. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 33(2), 263-282.
- Cooper, R. (2005). *Classifying madness* (Vol. 86). Dordrecht, The Netherlands: Springer.
- Cummins, R. (1975), "Functional Analysis", *The Journal of Philosophy* 72: 741-765. [Excerpt reprinted in N. Block (ed.), (1980), *Readings in Philosophy of Psychology*. Cambridge, MA: Harvard University Press, pp. 185-190.]
- Currie, A. M. (2014). Narratives, mechanisms and progress in historical science. *Synthese*, 191(6), 1163-1183.
- Darwin, C. (1872). *On the origin of species*. 6th Edition, online: <http://darwin-online.org.uk/content/frameset?itemID=F391&viewtype=text&pageseq=1>
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Dawkins, R. (1983). Adaptationism was always predictive and needed no defense. *Behavioral and Brain Sciences*, 6(3), 360-361.
- Dawkins, R. (1986). The Blind Watch maker. *Longman*, 1, 986.
- Dawkins, R. (1996a). Human chauvinism. *Evolution*, 51(3), 1015-1021
- Dawkins, R. (1996b). *The blind watchmaker: Why the evidence of evolution reveals a universe without design*. WW Norton & Company.
- Dawkins, R. (1997). *Climbing mount improbable*. WW Norton & Company.
- Dawkins, R. (2003). The evolution of evolvability. *On growth, form and computers*, 239-255.
- Dawkins, R. (2004). Extended phenotype—but not too extended. A reply to Laland, Turner and Jablonka. *Biology and Philosophy*, 19(3), 377-396.
- Dawkins, R. (2005). The illusion of design. *Skeptic (Altadena, CA)*, 12(2), 51-54
- Dawkins, R. (2012). The descent of Edward Wilson. *Prospect Magazine*, 195, 1-59.
- Dawkins, R. & Krebs, J. R. (1979) “*Arms Races Between And Within Species*” *Pro R Soc Lond B* 205:55-77
- Dennett, D. C. (1990). Evolution, error, and intentionality. *Contemporary Materialism*, 254.
- Dennett, D. C. (1991). Real patterns. *The journal of Philosophy*, 88(1), 27-51.
- Dennett, D. C. (1995a) “*Darwin’s Dangerous Idea*”. Penguin Books.

- Dennett, D. C. (1995b). Cognitive science as reverse engineering several meanings of “Top-down” and “Bottom-up”. In *Studies in Logic and the Foundations of Mathematics* (Vol. 134, pp. 679-689). Elsevier.
- Dobzhansky, T. (1956). The Biological Basis of Human Freedom. Page Barbour Lectures for 1954 at the University of Virginia.
- Doolittle, W. F. (1994) Evolutionary creativity and complex adaptations: a molecular biologist’s perspective. In Campbell, J. H., & Schopf, J. W. (Eds.). (1994). *Creative Evolution?!*. Jones & Bartlett Learning.
- Doolittle, W. F. and Bapteste, E. (2007) "Pattern pluralism and the Tree of Life hypothesis." *Proceedings of the National Academy of Sciences* 104.7: 2043-2049.
- Embley, T. M., & Martin, W. (2006). Eukaryotic evolution, changes and challenges. *Nature*, 440(7084), 623.
- Emerson, B. C. (2002). Evolution on oceanic islands: molecular phylogenetic approaches to understanding pattern and process. *Molecular ecology*, 11(6), 951-966.
- Emerson, B. C., & Patiño, J. (2018). Anagenesis, cladogenesis, and speciation on islands. *Trends in ecology & evolution*, 33(7), 488-491.
- Endler, J. A. (1986). Natural selection in the wild (No. 21). *Princeton University Press*.
- Ferguson, K. G. (2007). Biological function and normativity. *Philo*, 10(1), 17-26.
- Fisher, D. C. (1986). Progress in organismal design. In *Patterns and Processes in the History of Life* (pp. 99-117). Springer, Berlin, Heidelberg.
- Fisher, R. A. (1954). Retrospect of the criticism of the theory of natural selection in Evolution as a Process. In Mayr, E., Huxley, J., Hardy, A. C., & Ford, E. B. (1954). *Evolution as a Process*. *Evolution as a process*. Allen and Unwin, London, 105.
- Fodor, J. (1997). Special sciences: Still autonomous after all these years. *Philosophical perspectives*, 11(s11), 149-63.
- Forber, P. (2005) On the explanatory roles of natural selection. *Biology and Philosophy* 20.2-3: 329-342.
- Forber, P. (2010). Confirmation and explaining how possible. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 41(1), 32-40.
- Fortey, R. (2001). The Cambrian explosion exploded?. *Science*, 293(5529), 438-439.
- Francis, R. C. (2005). *Why men won't ask for directions: the seductions of sociobiology*. Princeton University Press.
- Fuller, S. (2018) Progressive eugenics is still with us, it has just politically evolved. <https://www.acsh.org/news/2018/01/15/progressive-eugenics-still-us-it-has-just-politically-evolved-12417>
- Futuyma, D. J. (1987). On the role of species in anagenesis. *The American Naturalist*, 130(3), 465-473.
- Futuyma, D. J. (2015). Can modern evolutionary theory explain macroevolution?. In *Macroevolution* (pp. 29-85). Springer, Cham.
- Gardner, A., & Grafen, A. (2009). Capturing the superorganism: a formal theory of group adaptation. *Journal of evolutionary biology*, 22(4), 659-671.
- Ghiselin, M. T. (1999). Progress and the economy of nature. *Journal of Bioeconomics*, 1(1), 35-45.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of science*, 69(S3), S342-S353.
- Glennan, S. (2009) "Productivity, relevance and natural selection." *Biology & Philosophy* 24.3: 325-339.

- Glennan, S. (2010). Ephemeral mechanisms and historical explanation. *Erkenntnis*, 72(2), 251-266.
- Glymour, B. (2001). Selection, indeterminism, and evolutionary theory. *Philosophy of Science*, 68(4), 518-535.
- Godfrey-Smith, P. (1994). A Modern History Theory of Functions. *Noûs*, 28(3), 344-362.
- Godfrey-Smith, P. (1999). Adaptationism and the Power of Selection. *Biology and Philosophy*, 14(2), 181-194.
- Godfrey-Smith, P. (2001). Three kinds of adaptationism. *Adaptationism and optimality*, 122.
- Godfrey-Smith, P. (2007). Conditions for evolution by natural selection. *The Journal of Philosophy*, 104(10), 489-516.
- Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. Oxford University Press.
- Gould, S. J. (1988a). Trends as changes in variance: a new slant on progress and directionality in evolution. *Journal of Paleontology*, 62(3), 319-329.
- Gould, S. J. (1988b). *On replacing the idea of progress with an operational notion of directionality*. In Nitecki, M. H. (1988). *Evolutionary progress*. University of Chicago Press.
- Gould, S. J. (1990). *Wonderful life: the Burgess Shale and the nature of history*. WW Norton & Company.
- Gould, S. J. (1996) "Full House: The Spread of Excellence From Plato to Darwin" Harmony Books
- Gould, S. J. (1997a) "Self-Help For A Hedgehog Stuck On A Molehill: Review Of *Climbing Mount Improbable*" *Evolution* 51/3
- Gould, S. J. (1997b) "Darwinian Fundamentalism" Parts 1&2. The New York Review Of Books.
- Gould, S. J. (1997c). Redrafting the tree of life. *Proceedings of the American Philosophical Society*, 141(1), 30-54.
- Gould, S. J. (1998). Full house: the spread of excellence from Plato to Darwin. *Senior Managing Editor*, 5(2), 68.
- Gould, S. J. (2002a) "*Ontogeny And Phylogeny*". The Belknap Press Of The Harvard University Press.
- Gould, S. J. (2002b). Life's little Joke. *Resonance*, 7(11), 87-98.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 205(1161), 581-598.
- Gould, S. J., & Vrba, E. S. (1982). Exaptation—a missing term in the science of form. *Paleobiology*, 8(1), 4-15.
- Grafen, A. (1984). Natural selection, kin selection and group selection. *Behavioural ecology: An evolutionary approach*, 2, 62-84.
- Gray, M. W., Burger, G., & Lang, B. F. (1999). Mitochondrial evolution. *Science*, 283(5407), 1476-1481.
- Gregory, T. R. (2008). The evolution of complex organs. *Evolution: Education and Outreach*, 1(4), 358.
- Griffiths, P. E. (1996). The historical turn in the study of adaptation. *The British journal for the philosophy of science*, 47(4), 511-532.

- Hahlweg, K. (1991) "On The Notion Of Evolutionary Progress" *Philosophy Of Science* 58/3:436-451
- Hallam, A. (1997). Speciation patterns and trends in the fossil record. *Geobios*, 30(7), 921-930
- Hallam, A. (2009). *The problem of punctuational speciation and trends in the fossil record* (pp. 423-432). University of Chicago Press: Chicago, IL, USA.
- Hesketh, I. (2016). Counterfactuals and history: Contingency and convergence in histories of science and life. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 58, 41-48.
- Hoehn, K. B., Harnik, P. G., & Roth, V. L. (2016). A framework for detecting natural selection on traits above the species level. *Methods in Ecology and Evolution*, 7(3), 331-339.
- Hume, D. (1739). *A Treatise on Human Nature*. Clarendon Press.
- Huxley, J. (1942). *Evolution. The modern synthesis*. George Allen & Unwin Ltd.
- Jablonski, D. (1986). Background and mass extinctions: the alternation of macroevolutionary regimes. *Science*, 231(4734), 129-133.
- Jablonski, D., Gould, S. J., & Raup, D. M. (1986). The nature of the fossil record: a biological perspective. In *Patterns and processes in the history of life* (pp. 7-22). Springer, Berlin, Heidelberg.
- Jackson, F., & Pettit, P. (1992). In defense of explanatory ecumenism. *Economics & Philosophy*, 8(1), 1-21.
- Joslin, D. (2006). Real realization: Dennett's real patterns versus Putnam's ubiquitous automata. *Minds and Machines*, 16(1), 29-41.
- Joyce, R. (2007) "The Evolution Of Morality" M.I.T. Press
- Johnston, C. (2014) https://www.theguardian.com/science/2014/nov/07/richard-dawkins-labelled-journalist-by-eo-wilson?CMP=Share_iOSApp_Other
- Kimura, M. (1961). Natural selection as the process of accumulating genetic information in adaptive evolution. *Genetics Research*, 2(1), 127-140.
- Kimura, M. (1968). Evolutionary rate at the molecular level. *Nature*, 217(5129), 624-626.
- Kimura, M. (1994). Stochastic processes and distribution of gene frequencies under natural selection (1955). *Population Genetics, Molecular Evolution, and the Neutral Theory: Selected Papers*. The University of Chicago Press, Chicago, 144.
- Kingma, E. (2014, November). Naturalism about health and disease: adding nuance for progress. In *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine* (Vol. 39, No. 6, pp. 590-608). Journal of Medicine and Philosophy Inc..
- Kitcher, P. (1993). Function and design. *Midwest studies in philosophy*, 18(1), 379-397.
- Koonin, E. V. (2010). The origin and early evolution of eukaryotes in the light of phylogenomics. *Genome biology*, 11(5), 209.
- Ladyman, J., Ross, D., Collier, J., Spurrett, D., Spurrett, D., & Collier, J. G. (2007). *Every thing must go: Metaphysics naturalized*. Oxford University Press on Demand.
- Land, M. F., & Nilsson, D. E. (2002) *Animal eyes*. Oxford University Press.
- Lamarck, J. B. (1809). Zoological Philosophy, trans. *Hugh Elliot (New York: Hafner, 1963)*, 109-111.

- Langdon, J. H. (1997). Umbrella hypotheses and parsimony in human evolution: a critique of the Aquatic Ape Hypothesis. *Journal of Human Evolution*, 33(4), 479-494.
- Layne Jr, J. R., & Lee Jr, R. E. (1995). Adaptations of frogs to survive freezing. *Climate Research*, 5(1), 53-59.
- Levins, R., & Lewontin, R. C. (1985). *The dialectical biologist*. Harvard University Press.
- Lewens, T. (2004). Organisms and Artifacts Design in Nature and Elsewhere. *The MIT Press*.
- Lewens, T. (2009). Seven types of adaptationism. *Biology & Philosophy*, 24(2), 161.
- Lewontin, R. C. (1970). The units of selection. *Annual review of ecology and systematics*, 1(1), 1-18.
- Losos, J. B (2011) "Convergence, Adaptation And Constraint". *Evolution* 65/7: 1827-1840
- Losos, J. B., Jackman, T. R., Larson, A., de Queiroz, K., & Rodríguez-Schettino, L. (1998). Contingency and determinism in replicated adaptive radiations of island lizards. *Science*, 279(5359), 2115-2118.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of science*, 67(1), 1-25.
- Mackie, J.L., (1977). *Ethics: Inventing Right and Wrong*, New York: Penguin.
- Matthen, M., & Ariew, A. (2009). Selection and causation. *Philosophy of science*, 76(2), 201-224.
- Matthewson, J., & Calcott, B. (2011). Mechanistic models of population-level phenomena. *Biology & Philosophy*, 26(5), 737-756.
- Matthewson, J., & Calcott, B. (2017) Mechanistic Explanation without Mechanisms.
- Matthewson, J., & Griffiths, P. E. (2017). Biological criteria of disease: Four ways of going wrong. *Journal of Medicine and Philosophy*, 42(4), 447-466.
- Maynard-Smith, J. & Szathmary, E. (1995) "*The Major Transitions In Evolution*". W.H.Freeman Spektrum
- Mayr, E. (1982). The growth of biological thought: Diversity, evolution, and inheritance. *Harvard University Press*.
- Mayr, E. (1988). Toward a new philosophy of biology: Observations of an evolutionist (No. 211). *Harvard University Press*.
- McConwell, A. K., and Currie, A. (2017) "Gouldian arguments and the sources of contingency." *Biology & Philosophy* 32.2: 243-261.
- McShea, D. W. (1998). Possible largest-scale trends in organismal evolution: Eight "Live Hypotheses". *Annual review of ecology and systematics*, 29(1), 293-318.
- McShea, D. W. (2016). Three trends in the history of life: an evolutionary syndrome. *Evolutionary Biology*, 43(4), 531-542.
- McShea, D. W., & Brandon, R. N. (2010). *Biology's first law: the tendency for diversity and complexity to increase in evolutionary systems*. University of Chicago Press.
- McShea, D. W., & Simpson, C. (2011). The miscellaneous transitions in evolution. *The major transitions in evolution revisited*, 19-34.
- Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT press.

- Millikan, R. G. (1989a). An ambiguity in the notion "function". *Biology and Philosophy*, 4(2), 172-176.
- Millikan, R. G. (1989b) "In Defense Of Proper Functions" *Philosophy Of Science* 56/2: 288-302
- Millstein, R. L. (2000) "Chance and macroevolution." *Philosophy of Science* 67.4: 603-624.
- Millstein, R. L. (2006). Natural selection as a population-level causal process. *The British Journal for the Philosophy of Science*, 57(4), 627-653.
- Moore, G. E. (1959). *Principia Ethica*:(1903) (Vol. 2). *Cambridge University Press*
- Morgan, E. (1982). *The aquatic ape*. Scarborough House.
- Morris, S. C. (2003). *Life's solution: inevitable humans in a lonely universe*. Cambridge University Press.
- Moss, L. (2012). Is the philosophy of mechanism philosophy enough?. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(1), 164-172.
- Muller, G. B., & Wagner, G. P. (1991). Novelty in evolution: restructuring the concept. *Annual review of ecology and systematics*, 22(1), 229-256.
- Nanay, B. (2005) "Can cumulative selection explain adaptation?." *Philosophy of Science* 72.5: 1099-1112.
- Neander, K. (1988) What does natural selection explain? Correction to Sober. *Philosophy of Science* 55.3: 422-426.
- Neander, K. (1991a) "The teleological notion of function" *Australasian journal of philosophy*. 69/4: 454-468
- Neander, K. (1991b). Functions as selected effects: The conceptual analyst's defense. *Philosophy of science*, 58(2), 168-184.
- Neander, K. (1995a). Pruning the tree of life. *The British Journal for the Philosophy of Science*, 46(1), 59-80.
- Neander, K. (1995b). Explaining complex adaptations: A reply to Sober's' reply to Neander'. *The British journal for the philosophy of science*, 46(4), 583-587.
- Neander, K. (1995c). Misrepresenting & malfunctioning. *Philosophical Studies*, 79(2), 109-141.
- Nichols, S. (2004). *Sentimental rules: On the natural foundations of moral judgment*. *Oxford University Press*.
- Nilsson, D. E., & Pelger, S. (1994). A pessimistic estimate of the time required for an eye to evolve. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 256(1345), 53-58.
- Nicholson, D. J. (2012). The concept of mechanism in biology. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(1), 152-163.
- Nitecki, M. H. (Ed) (1988) "Evolutionary Progress". The University Of Chicago Press
- Nowak, M. A., & Sigmund, K. (2004). Evolutionary dynamics of biological games. *science*, 303(5659), 793-799.
- O'Hara, R. J. (1992) "Telling the tree: narrative representation and the study of evolutionary history." *Biology and Philosophy* 7.2: 135-160.
- Okasha, S. (2006a) "The Levels Of Selection Debate: Philosophical Issues". *Philosophy Compass* 1/1: 74-85
- Okasha, S. (2006b) *Evolution And The Levels Of Selection*. *Oxford University Press*

- Okasha, S. (2008) "Fisher's Fundamental Theorem Of Natural Selection- A Philosophical Analysis" *Brit. J. Phil. Sci* 0:1-33
- Okasha, S. "Reply To Sober And Waters" (2010) *Philosophy And Phenomenological Research* 82/1: 241-248
- Okasha, S. (2012a) "Wynne-Edwards and the history of group selection" *Metascience* 21(2), 355-357
- Okasha, S. (2012b) "Population Genetics" *Stanford Encyclopedia, Online.*
- Okasha, S. (2018a) "The strategy of endogenization in evolutionary biology." *Synthese*: 1-23.
- Okasha, S. (2018b). *Agents and goals in evolution.* Oxford University Press.
- Okasha, S, and Paternotte, C. (2012) "Group adaptation, formal Darwinism and contextual analysis." *Journal of evolutionary biology* 25.6: 1127-1139.
- O'Hara, R. J. (1992). Telling the tree: narrative representation and the study of evolutionary history. *Biology and Philosophy*, 7(2), 135-160.
- O'Malley, M. A. (2010). The first eukaryote cell: an unfinished history of contestation. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 41(3), 212-224.
- O'Malley, M. A., & Powell, R. (2016). Major problems in evolutionary transitions: how a metabolic perspective can enrich our understanding of macroevolution. *Biology & Philosophy*, 31(2), 159-189.
- Orr, H. A. (2005). The probability of parallel evolution. *Evolution*, 59(1), 216-220.
- Orzack, S. H., & Sober, E. (1994). Optimality models and the test of adaptationism. *The American Naturalist*, 143(3), 361-380.
- Orzack, S. H., & Sober, E. (1996). How to formulate and test adaptationism. *The American Naturalist*, 148(1), 202-210.
- Otto, S. P., Servedio, M. R., & Nuismer, S. L. (2008). Frequency-dependent selection and the evolution of assortative mating. *Genetics*, 179(4), 2091-2112.
- Papineau, D. (1984). Representation and explanation. *Philosophy of Science*, 51(4), 550-572.
- Parker, G. A. (1983). Arms races in evolution—an ESS to the opponent-independent costs game. *Journal of Theoretical Biology*, 101(4), 619-648.
- Parker, G. A., & Smith, J. M. (1990). Optimality theory in evolutionary biology. *Nature*, 348(6296), 27.
- Pettitt, J. M., McConchie, C. A., Ducker, S. C., & Knox, R. B. (1980). Unique adaptations for submarine pollination in seagrasses. *Nature*, 286(5772), 487.
- Powell, R. (2012). Convergent evolution and the limits of natural selection. *European Journal for Philosophy of Science*, 2(3), 355-373.
- Priest, G. (2017). Framing causal questions about the past: The Cambrian explosion as case study. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 63, 55-63.
- Quandt, E. M., Deatherage, D. E., Ellington, A. D., Georgiou, G., & Barrick, J. E. (2014). Recursive genomewide recombination and sequencing reveals a key refinement step in the evolution of a metabolic innovation in *Escherichia coli*. *Proceedings of the National Academy of Sciences*, 111(6), 2217-2222.
- Railton, P., (1986) "Moral Realism," *Philosophical Review*, 4(3): 163–207.

- Raup, D. M., & Jablonski, D. (1986). *Patterns and processes in the history of life: report of the Dahlem Workshop on Patterns and Processes in the History of Life, Berlin 1985, June 16-21*. Springer-Verlag.
- Razeto-Barry, P. (2013). Complexity, adaptive complexity and the Creative View of natural selection. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(3), 312-315.
- Razeto-Barry, P., & Frick, R. (2011). Probabilistic causation and the explanatory role of natural selection. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 42(3), 344-355.
- Rensch, B. (1959). *Evolution above the species level*. Columbia University press.
- Reznick, D. (1985). Costs of reproduction: an evaluation of the empirical evidence. *Oikos*, 257-267.
- Richards, R. J. (1992). The meaning of evolution: The morphological construction and ideological reconstruction of Darwin's theory. *University of Chicago Press*
- Ridley, M. (2007). *Evolution*. Blackwell Publishing
- Rosenberg, A. (1994). Instrumental biology, or the disunity of science. *University of Chicago Press*.
- Ross, D. (2000). Rainforest realism: A Dennettian theory of existence. *Dennett's philosophy: A comprehensive assessment*, 147-168.
- Rosslénbroich, B. (2009). The theory of increasing autonomy in evolution: a proposal for understanding macroevolutionary innovations. *Biology & Philosophy*, 24(5), 623-644.
- Rosslénbroich, B. (2014). *On the origin of autonomy: a new look at the major transitions in evolution* (Vol. 5). Springer Science & Business Media.
- Ruse, M. (1993) "Evolution And Progress" *Tree* 8/2:55-59
- Ruse, M. (1995). *Evolutionary naturalism: selected essays*. Routledge.
- Ruse, M. (2006). "Richard Dawkins and the problem of progress" in Grafen, A., & Ridley, M. (Eds.). (2007). *Richard Dawkins: how a scientist changed the way we think: reflections by scientists, writers, and philosophers*. Oxford University Press, USA.
- Ruse, M. (2009). *Monad to man: the concept of progress in evolutionary biology*. Harvard University Press.
- Runnegar, B. (1982). The Cambrian explosion: animals or fossils?. *Journal of the Geological Society of Australia*, 29(3-4), 395-411.
- Sansom, R. (2003). Constraining the adaptationism debate. *Biology and Philosophy*, 18(4), 493-512.
- Seager, W. (2000). Real patterns and surface metaphysics. *Dennett's Philosophy: A Comprehensive Assessment*, 95-130.
- Shanahan, T. (2004). *The evolution of Darwinism: selection, adaptation and progress in evolutionary biology*. Cambridge University Press.
- Simpson, C. (2010). Species selection and driven mechanisms jointly generate a large-scale morphological trend in monobathrid crinoids. *Paleobiology*, 36(3), 481-496.
- Simpson, G. G. (1949). *The meaning of evolution: a study of the history of life and of its significance for man* (Vol. 25). Yale University Press.
- Simpson, G. G. (1974). The concept of progress in organic evolution. *Social Research*, 28-51.

- Skipper Jr, R. A., & Millstein, R. L. (2005). Thinking about evolutionary mechanisms: Natural selection. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), 327-347.
- Smith, J. M., & Szathmáry, E. (1997). *The major transitions in evolution*. Oxford University Press.
- Sober, E. (1984). *The nature of selection: Evolutionary theory in philosophical focus*. University of Chicago Press.
- Sober, E. (1992). Screening-off and the units of selection. *Philosophy of science*, 59(1), 142-152.
- Sober, E. (1993). *The nature of selection: Evolutionary theory in philosophical focus*. University of Chicago Press.
- Sober, E. (1995). Natural selection and distributive explanation: A reply to Neander. *The British journal for the philosophy of science*, 46(3), 384-397.
- Sober, E., & Orzack, S. H. (2003). Common ancestry and natural selection. *The British journal for the philosophy of science*, 54(3), 423-437.
- Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge University Press.
- Spencer, H. (1891) *Progress: Its Law And Cause*. https://openlibrary.org/books/OL24183555M/Progress_its_law_and_cause
- Stearns, S. C. (2000). Life history evolution: successes, limitations, and prospects. *Naturwissenschaften*, 87(11), 476-486.
- Stegmann, U. E. (2010). What can natural selection explain?. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 41(1), 61-66.
- Sterelny, K. (1996). Explanatory pluralism in evolutionary biology. *Biology and Philosophy*, 11(2), 193-214.
- Sterelny, K. (1999) "Bacteria At The High Table" *Biology And Philosophy* 14:459-470
- Sterelny, K. (2004). Symbiosis, evolvability and modularity. *Modularity in development and evolution*, 490-516.
- Sterelny, K., & Fraser, B. (2016). Evolution and moral realism. *The British Journal for the Philosophy of Science*, 68(4), 981-1006.
- Sterelny, K., & Griffiths, P. E. (2012). *Sex and death: An introduction to philosophy of biology*. University of Chicago press.
- Sterelny, K., & Kitcher, P. (1988). The return of the gene. *The Journal of Philosophy*, 85(7), 339-361
- Szathmáry, E. (2015). Toward major evolutionary transitions theory 2.0. *Proceedings of the National Academy of Sciences*, 112(33), 10104-10111.
- Thompson, J. N. (1986). Constraints on arms races in coevolution. *Trends in ecology & evolution*, 1(4), 105-107.
- Tolliver, J. T. (1994). Interior colors. *Philosophical Topics*, 22(1/2), 411-441.
- Turelli, M., Barton, N. H., & Coyne, J. A. (2001). Theory and speciation. *Trends in ecology & evolution*, 16(7), 330-343.
- Van Valen, L. (1977). The red queen. *The American Naturalist*, 111(980), 809-810.

- Varki, A. (2016). Why are there no persisting hybrids of humans with Denisovans, Neanderthals, or anyone else?. *Proceedings of the National Academy of Sciences*, 113(17), E2354-E2354.
- Vavova, K. (2015). Evolutionary debunking of moral realism. *Philosophy Compass*, 10(2), 104-116.
- Vermeij, G. J. (1999). Inequality and the directionality of history. *The American Naturalist*, 153(3), 243-253.
- Vermeij, G. J. (2006) "Historical Contingency And The Purported Uniqueness Of Evolutionary Innovations" PNAS 103/6: 1804-1809
- Vorzimmer, P. (1963). Charles Darwin and blending inheritance. *Isis*, 54(3), 371-390.
- Wake, D. B., Wake, M. H., & Specht, C. D. (2011). Homoplasy: from detecting pattern to determining process and mechanism of evolution. *science*, 331(6020), 1032-1035.
- Walsh, D. M. (1998). The scope of selection: Sober and Neander on what natural selection explains." *Australasian Journal of Philosophy* 76.2: 250-264.
- Walsh, D. M. (2000). Chasing shadows: natural selection and adaptation. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 31(1), 135-153.
- Walsh, D. M., Lewens, T., & Ariew, A. (2002). The trials of life: Natural selection and random drift. *Philosophy of Science*, 69(3), 429-446
- Wilkerson, W. S. (1997). Real patterns and real problems: Making Dennett respectable on patterns and beliefs. *The Southern journal of philosophy*, 35(4), 557-570.
- Wilson, D. S. (2012). Richard Dawkins, Edward O. Wilson, and the consensus of the many. *This view of life*, 29.
- Wilson, D. S., & Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and brain sciences*, 17(4), 585-608.
- Wilson, E. O. (1991). *The Diversity Of Life*. Penguin Books.
- Wimsatt, W. C. (2001). Generative entrenchment and the developmental systems approach to evolutionary processes. *Cycles of contingency: Developmental systems and evolution*, 219-237.
- Woodward, J. (2002). What is a mechanism? A counterfactual account. *Philosophy of science*, 69(S3), S366-S377.
- Woodward, J. (2002). What is a mechanism? A counterfactual account. *Philosophy of science*, 69(S3), S366-S377.
- Wouters, A. G. (2003). Four notions of biological function. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 34(4), 633-668.
- Wray, G. A., Levinton, J. S., & Shapiro, L. H. (1996). Molecular evidence for deep Precambrian divergences among metazoan phyla. *Science*, 274(5287), 568-573.
- Wright, L. (1973). Function. *Philosophical Review*, 82, 139-168.
- Yule, J., Fournier, R., & Hindmarsh, P. (2013). Biodiversity, extinction, and humanity's future: The ecological and evolutionary consequences of human population and resource use. *Humanities*, 2(2), 147-159.
- Yoshimura, J., & Shields, W. M. (1987). Probabilistic optimization of phenotype distributions: a general solution for the effects of uncertainty on natural selection?. *Evolutionary Ecology*, 1(2), 125-138.

