Reading the Book of Life:

Contingency and Convergence in Macroevolution

by

Russell Powell

Department of Philosophy
Duke University

Date:_____

Approved:

_____
Alexander Rosenberg, Supervisor

_____
Robert N. Brandon

_____
Owen J. Flanagan

_____
Dan W. McShea

_____
Karen Neander

_____
V. Louise Roth

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Philosophy in the Graduate School
of Duke University

2008

ABSTRACT

Reading the Book of Life:

Contingency and Convergence in Macroevolution

by

Russell Powell

Department of Philosophy
Duke University

Date:_____

Approved:

_____
Alexander Rosenberg, Supervisor

_____
Robert N. Brandon

_____
Owen J. Flanagan

_____
Dan W. McShea

_____
Karen Neander

_____
V. Louise Roth

An abstract of a dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Philosophy in the Graduate School
of Duke University

2008

ABSTRACT

This dissertation explores philosophical problems in biology, particularly those relating to macroevolutionary theory. It is comprised of a series of three papers drawn from work that is currently at the publication, re-submission, and review stage of the journal refereeing process, respectively. The first two chapters concern the overarching contours of complex life, while the third zeroes in on the short and long-term prospects of human evolution.

The rhetorical journey begins with a thought experiment proposed by the late paleontologist Stephen Jay Gould. Gould hypothesized that replaying the "tape of life" would result in radically different evolutionary outcomes, both with respect to animal life in general and the human species in particular. Increasingly, however, biologists and philosophers are pointing to convergent evolution as evidence for replicability and predictability in macroevolution. Chapters 1 and 2 are dedicated to fleshing out the Gouldian view of life and its antithesis, clarifying core concepts of the debate (including contingency, convergence, constraint and causation), and interpreting the empirical data in light of these conceptual clarifications. Chapter 3 examines the evolutionary biological future of the human species, and the ways in which powerful new biotechnologies can

shape it, for better and for worse. More detailed chapter summaries are provided below.

In **Chapter 1**, I critique a book-length excoriation of Gould's contingency theory written by the paleobiologist Simon Conway Morris, in which he amasses and marshals a good bulk of the homoplasy literature in the service of promoting a more robust, counter-factually stable account of macroevolution. I show that there are serious conceptual and empirical difficulties that arise in broadly appealing to the frequency of homoplasy as evidence for robustness in the history of life. Most important is Conway Morris's failure to distinguish between convergent ('externally' constrained) and parallel ('internally' constrained) evolution, and to consider the respective implications of these significantly different sources of homoplasy for a strong adaptationist view of life.

In so doing, I propose a new definition of parallel evolution, one intended to rebut the common charge that parallelism differs from convergence merely in degree and not in kind. I argue that although organisms sharing a homoplastic trait will also share varying degrees of homology (given common decent), it is the underlying developmental homology with respect to the generators directly causally responsible for the homoplastic event that defines parallel evolution and non-arbitrarily distinguishes it from convergence. I make use of the philosophical concept

of 'screening-off' in order to distinguish the proximate generators of a homoplastic trait from its more distal genetic causes (such as conserved master control genes).

In **Chapter 2**, I critically examine a recent assessment of the contingency debate by the philosopher John Beatty, in which he offers an interpretation of Gould's thesis and argues that it is undermined by iterative ecomorphological evolution. I develop and defend alternative concepts of contingency and convergence, and show how much of the evidence generally held to negate the contingency thesis not only fails to do so, but in fact militates in favor of the Gouldian view of life. My argument once again rests heavily on the distinction between parallelism and convergence, which I elaborate on and defend against a recent assault by developmental biologists, in part by recourse to philosophical work on the ontological prioritization of biological causes.

In **Chapter 3**, I explore the probable (and improbable) evolutionary biological consequences of intentional germ-line modification, particularly in relation to human beings. A common worry about genetic engineering is that it will reduce the pool of genetic diversity, creating a biological monoculture that could not only increase our susceptibility to disease, but even hasten the extinction of our species. Thus far, however, the evolutionary implications of human genetic modification have remained

largely unexplored. In this Chapter, I consider whether the widespread use of genetic engineering technology is likely to narrow the present range of genetic variation, and if so, whether this would in fact lead to the evolutionary harms that some authors envision. By examining the nature of biological variation and its relation to population immunity and evolvability, I show that not only will genetic engineering have a negligible impact on human genetic diversity, but that it will be more likely to ensure rather than undermine the health and longevity of the human species. To this end, I analyze the relationship between genotypic and phenotypic variation, consider process asymmetries between micro and macroevolution, and investigate the relevance of evolvability to clade-level persistence and extinction.

Key words: Constraint, Contingency, Convergence, Evolvability, Genetic Engineering, Homoplasy, Macroevolution, Parallelism, Variation

***

CONTENTS

Chapter 3

The Evolutionary Biological Implications of Human Genetic Engineering

LIST OF FIGURES

Chapter 1

Is Convergence More Than an Analogy? Homoplasy and the Nature of Macroevolution

1. Introduction

It is widely accepted among biologists and philosophers of biology that replaying the proverbial 'tape of life' would result in wildly unpredictable and radically different evolutionary outcomes. Some authors even espouse the more radical notion that virtually every interesting event in the history of life falls into the realm of historical contingency. This position arises from the empirical assumption that the invariant laws of nature carve out exceedingly broad channels that only loosely constrain the evolution of organismic design. On this view, observed evolutionary products are sensitively dependent on stochastic initial conditions, a feature of macroevolutionary processes which undermines attempts to formulate robust generalizations regarding the evolution of organismic form. Because even the most successful biological generalizations are undermined by multiply realizable solutions to contingent design problems (Beatty 1995) that are destined for obsolescence in the unrelenting arms race of natural selection (Rosenberg 2001; Van Valen 1973), they tend to lack the nomic

necessity and counterfactual stability canonically characteristic of natural laws.

With this nomological vacuum as the backdrop, historicists tend to explain the inhomogenous distribution of organismic form in morphospace (at all levels of the genealogical hierarchy) as the result of internal developmental constraints that restrict the set of possible variants on which selection can operate. Strong adaptationists, on the other hand, tend to view the clumping of actualized morphology not as a contingent consequence of developmental inertia, but rather as an ecologically optimal set of solutions to relatively stable functional problems. Although historical, developmental, stochastic, and adaptive explanations of evolution are not exclusive of or even necessarily opposed to one another (Gould 2002; Sterelny 1996), counterfactual biologists may seek to individuate characters of taxa or even entire clades that are 'robust,' or insensitive to marginal historical perturbations.

A number of authors have touted examples of 'convergent evolution' as evidence for not only adaptation, but also for *hard adaptationism* (*sensu* Amundson 1994), or the view that nearly all scientifically interesting features of life, from speciel morphology to clade geometry, can be explained by natural selection.[1] Hard adaptationism assumes that design space is *highly constrained externally*, that is by directional and

[1] See e.g. Conway Morris (2003); Foley (1999); Dennett (1995).

subsequent stabilizing selection, the latter containing genetic drift once a (local) optimality has been achieved. However, it also presupposes that morphospace is virtually *unconstrained internally*, namely by postulating the inexorable tendency of natural selection to overcome developmental constraints that would otherwise lead to adaptive sub-optimality. Hard adaptationism is not philosophically opposed to contingency per se, but rather to the notion that 'frozen accidents'—or combinations of contingency plus developmental constraint—are of significant macroevolutionary consequence.

Of the self-proclaimed or justly labeled hard adaptationists, Simon Conway Morris ("SCM") is the most prominent champion of homoplasy and its purported implications for a non-contingent, counterfactually stable view of life. In essence, his view is this: The ubiquity of 'convergence,' or the independent origination of similar forms in distantly related organisms, is prima facie evidence for an adaptational design space that is so severely constrained by the invariant chemico-physical laws that organismic form will, through the optimizing forces of natural selection, repeatedly and inevitably converge on certain identifiable functional attractors or biological properties. According to SCM, homoplasy not only represents an important data set in the debate over frozen contingency, but bespeaks a deeper regularity underlying macrobiological pattern.

However, there are numerous conceptual and empirical problems that arise in broadly appealing to the frequency of homoplasy as evidence for a non-contingently constrained adaptational design space. Most important is the need to distinguish between convergent evolution (due to external, non-contingent constraint) and parallel evolution (due in large part to internal, contingent constraint), and to consider how the respective frequencies of these significantly different sources of homoplasy affect a strong adaptationist view of life. In this chapter, I critically evaluate SCM's use of the homoplasy literature in his attempt to bolster a functionalist account of macroevolution. In so doing, I offer a conception of parallelism which avoids the charge that it differs from convergence merely in degree and not in kind. I argue that although organisms sharing a homoplastic trait will also share varying degrees of homology, it is the underlying developmental homology *with respect to the generators directly causally responsible for the homoplastic event* that defines parallel evolution and non-arbitrarily distinguishes it from convergence.

2. Pervasive Homoplasy and Evolutionary Inevitability

Biologists and philosophers have described the evolutionary phenomenon of convergence as nature's way of re-winding the tape of life (Dennett 1995), biology's closest analog to independent experimental replication (Gould 2002). As one preeminent comparative physiologist suggests (Vogel 1996,

1998), the project of identifying convergence offers more than just evidence for adaptation, for it enables biologists to distinguish aspects of form that are strongly determined by functional demands from those that are less fundamental to design. In *The Crucible of Creation: The Burgess Shale and the Rise of Animals* (1998) and more recently in *Life's Solution: Inevitable Humans in a Lonely Universe* (2003), SCM sketches a view of life that is somewhat unusual from a contemporary philosophical and biological standpoint. Pointing to the putative pervasiveness of convergence, SCM defends the strong adaptationist view that life inexorably navigates toward certain pre-ordained "endpoints," to which "the routes are many, but the destinations few, and the landscapes across which all organisms must travel are adaptive" (2003, 297). Compiling a litany of classical and lesser known examples of homoplasy at all levels of the biological hierarchy, he concludes that contingency is not an important feature of macroevolution. For instance, after discussing the remarkable case of 'convergence' between the marsupial fauna of the super-islands of Australia and South America and the placental mammals, SCM goes on to render the following strong conclusion (1998, 205): "Does it follow then that contingent processes are an irrelevance in the way we see the world? I have argued that, so far as the history of life is concerned, they are."

In addition, SCM maintains that unlike the specific adaptations of particular species, convergence on certain "biological properties" suggests that they are facets of a robust evolutionary process that will, despite the non-linearity of their actual sequence, inevitably manifest at some spatio-temporal point in the unfolding of deep evolutionary time. He concludes that (2003, 308):

> [A]lthough any history is necessarily unique, the resultant complex end-form is not simply the contingent upshot of local and effectively random processes. On any other suitable planet there will I suggest be animals very much like mammals, and mammals much like apes. Not identical, but surprisingly similar.

In essence, then, SCM's challenge is not directed at uncontroversial theories of contingency associated with "the destiny of a given lineage" (1998, 201), but toward the more radical notion that most or all interesting biological properties are themselves highly sensitively dependent on initial conditions, including class-level properties like 'mammalness,' or even phyla-level properties such as 'arthropodness.' Instead, SCM asserts that natural selection is not only necessary but *sufficient* to explain lumpy morphospace occupation at all taxonomic levels.

Similarly enamored of convergence, Dennett (1995, 306), in referring to the Cambrian 'experiment' (the *locus classicus* of debates over contingency thanks to Gould 1989), contends that "*whichever* lineage

happens to survive will gravitate toward the Good Moves in Design Space" (emphasis in original). "Replay the tape a thousand times," Dennett claims, "and the Good Tricks will be found again and again" (308). What Dennett means by a "Good Trick" is not entirely clear, although given the context (and the caps) he must mean more than simply an adaptation; for to counter radical contingency, a Good Trick must entail optimality in a more global sense. SCM agrees with Dennett's position that "convergence…is the fatal weakness in [the] case for contingency" (Ibid). He holds that if only biologists could identify these elusive 'laws of convergence' operating beneath the surface of stochastic chaos, they would be able to divine robust, counterfactually resilient predictions regarding the evolution of life on Earth and throughout the Universe.

Pit against this radical functionalist *weltanschauung* is the 'historicist' (*sensu* Gould 2002) view of the history of life, which views the inhomogeneous distribution or 'clumping' of organisms across morphospace not as an optimal set of solutions to functional problems (courtesy of natural selection and the invariant physical laws), but as the result of internal, contingent constraints restricting the realm of the possible. Rather than a predictable march of organismic form toward identifiable optimality or equilibrium, the historicist's history of life is (as Henry Ford was fond of saying) simply 'one damn thing after another'—an

accumulation of accidents, intelligible in hindsight but wildly unpredictable in prospect, dancing in rhythm with a stochastic ecology and exhibiting no long-term bias toward any particular functional solution. Of course, historicists agree with adaptationists that natural selection is the only known mechanism for producing function, but they deny the strong functionalist claim that selection is the predominant force behind macroevolutionary pattern.

Although historicists will tend to attribute macrobiological features (such as clade topography) to stochastic rather than competitive models of interaction, the crux of the historicist dispute with SCM is not the relative significance of selection per se, but rather that the former attributes the bulk of macroevolutionary change—whether functionally or stochastically driven—to contingent rather than nomologically inevitable events. Thus, contrary to popular conception, the contingency debate in biology is not between determinism and chance and their respective roles in evolution—for a chaotic world in which outcomes are hypersensitive to boundary conditions may be perfectly deterministic. Rather, the key question concerns whether the homoplastic evolutionary features identified by SCM reflect a nomic necessity that giant gold cubes and other accidents will never enjoy.

SCM for his part declares that contingent and stochastic processes play a *de minimus* role in shaping macroevolution over deep geological time, stating that such mechanisms are completely "irrelevant…in so far as the history of life is concerned" (1998, 205). This is a strong and iconoclastic claim, but in fact his assertions are even more radical, as he refers to the notion that contingent forces drive evolutionary processes as a biological fundamentalist "myth" (2003, 322). These views alone provide sufficient grounds for classifying SCM as a hard adaptationist, as they entail that (virtually) all scientifically significant features of life are of adaptive provenance, with selection overpowering any developmental constraints or stochastic tendencies (such as genetic drift).

Constraint in evolutionary biology may be viewed as circumstances limiting the nature of design problems and their set of possible solutions. External constraint is *non-contingent*, imposed by the chemico-physical laws and their interaction with the optimizing agency of natural selection. It is this constraint which could theoretically restrict the universe of functional solutions to a manageable handful that could admit of prediction without a burdensome litany of *ceteris paribus* qualifications. For instance, because of the laws of optics and the properties of light, there may be only a handful of ways to macroscopically arrange an image-forming organ of relatively high acuity—hence its convergent evolution up

to 15 times within and between distant phyla (Land 1992). Internal constraint, on the other hand, probably reflects not invariance but *frozen contingency*, or the radical conservation (via entrenchment and stabilizing selection) of upstream regulatory networks and even more distal sub-circuits that arose in response to local, stochastically fluctuating ecological pressures.[2]

In considering whether patterns of homoplasy are evidence for counterfactually stable limitations on organismic form, it is essential to recognize that each type of constraint could alone or jointly be responsible for the observed morphological regularities. As we shall see, SCM systematically conflates internal and external constraints on design space in arguing for a fundamentally non-contingent view of life.

3. Convergent versus Parallel Evolution

The most effective way of categorizing independently evolved similarities so as to reflect the above distinction between internal and external constraint is by recognizing two causally differentiated sub-categories within the larger rubric of homoplasy. These are *convergent* and *parallel* evolution. To clear up perceived confusion in the literature regarding the

[2] Gene regulatory networks are hierarchical, with earlier linkages having more pleiotropic effects than the more distal fine-grained terminal processes. The former upstream sub-circuits, which Davidson and Erwin (2006) have termed 'kernels,' specify the more general domain of the developing organism. Because kernels are 'recursively' wired, interference with any single kernel gene will destroy its function altogether, resulting in phenotypic catastrophe.

contrast between convergence and parallelism, Haas and Simpson (1946) published a seminal review of the topic. They defined *homology* as similarity based on common ancestry, in contradistinction to *homoplasy*, which simply refers to similarity that is *not* due to common descent. Unlike Haas who adopted a geometrical approach to convergence and parallelism[3], Simpson advocated a causal differentiation of the two concepts, with convergence resulting exclusively from common selective pressures, and parallelism linked to underlying homologous developmental pathways, or what he termed a "community of common ancestry" (1961, 103). The term *analogy* is often used to imply a common selected function as the underlying basis for a perceived similarity, whereas the broader term 'homoplasy' remains causally agnostic. Where analogy ultimately falls in the evolutionary lexicon, however, is not terribly relevant for the present purposes. What matters is whether convergence is more than an analogy in the *philosophical* (rather than technical) sense: That is, whether the existence of distantly related evolutionary simulacra indicates a deeper, predictable structure to macroevolution. Of greatest conceptual value to

[3] Haas preferred to distinguish convergence and parallelism geometrically (rather than causally)—with the former entailing that two lineages resemble one another more than their ancestors did, and the latter referring to cases in which two lineages evolve in the same direction without resembling one another any more than their ancestors did. For example, a trend of increasing body size in grasshopper and walrus clades would represent a parallelism for Haas, although Simpson would presumably demur since the parallel increases are probably not linked to shared developmental homology.

this investigation will be the contraposition of convergence with parallel evolution, each a type of homoplasy but with importantly different causal origins.

Whereas the distinction between homology and convergence is relatively crisp, the concept of parallelism occupies a gray zone between definitional homology (retention by common descent) and true convergence (similar design and function with entirely different developmental-structural origins) (Patterson 1988). Unlike convergence, parallelism contains too much 'homology-ness' to be considered *solely* the result of similar ecological pressures—that is to say, natural selection is necessary but not sufficient to explain parallelism. Since all extant organisms descend from a single common ancestor, they share important homologues (such as the general structure of their nucleic acids); recognizing this, some authors have concluded that the distinction between parallel and convergent evolution is ultimately one of degree rather than kind (Diogo 2005; Conway Morris 2003, 435 n.1). Nevertheless, although organisms sharing a homoplastic trait will also share varying degrees of homology, the decision to categorize a homoplastic event as one or the other is not arbitrary. This is because it is the underlying homology *with respect to the generators directly causally responsible for the homoplastic event* that

defines parallelism and offers a principled basis on which to distinguish it from convergence.

For instance, while Pax-6/eyeless (the poster boy of 'deep homology') may be an ancient master control gene conserved between metazoan phyla that is integral to the development of both vertebrate and mollusk camera-type eyes (Zuker 1994), it is not directly causally responsible for their analogous *macroscopic arrangements*, which are produced by wholly different developmental generators and processes and thus represent convergent rather than parallel evolution. While Gould (2002, 1159) may be correct to point out that deep homologues imply a more prominent role for parallel evolution than anticipated by the modern synthesis, the presence of a conserved master control gene does not make the evolution of any adaptation arising from generators subsequent in the developmental cascade a parallel event. To the contrary, we should meaningfully label as parallelism only cases in which an iterative morphology is proximally associated with a single regulatory gene or suite of genes that is conserved throughout a clade, merely to be 'switched' on and off by natural selection in accordance with the dictates of local ecology. Thus, the issue is not whether similar phenotypes share developmental origins at *any* phylogenetic depth, but rather whether the homoplastic trait in question

derives immediately from a shared set of developmental generators at the phylogenetic (and/or ontogenetic) depth relevant for that particular trait.

Deep homology may smack of contingency and frozen accident, but *contra* Gould (2002) such radical conservation is not the basis of parallelism if numerous processes along the developmental cascade intervene to proximately produce the homoplastic trait. Brandon's (1990) notion of 'screening-off' (adapted from Salmon 1971) may be helpful here in considering how the proximate developmental cause of a trait should be delineated. Proximal genetic cause P screens-off more distal cause D (e.g. a shared master control gene) of homoplastic trait T where the probability of T given P and D, is the same as the probability of T given P, and different from the probability of T given D (more formally, P screens off D from T iff $Pr(T,P\&D) = Pr(T,P) \neq Pr(T,D)$). The probability of T given P need not and will rarely be close to 1.0, as development is inherently noisy, affected by non-genetic conditions, and probably susceptible to quantum effects. But probabilistic unity is not required by the notion of screening-off, which is simply an asymmetric two-place causal relation, in this case between D and P with respect to T. This contingent (empirical) relation holds for master control genes like Pax-6/eyeless, which despite their central developmental role, do not even come close to exhausting the causal factors relevant to the production of cephalopod and vertebrate camera-type eyes,

and thus cannot serve as the basis on which to describe this homoplastic event as a parallelism.

Two basic experimental manipulations (shown in **Figure 1**) can be carried out in order to test whether P screens-off D in the case of Pax-6 (or any similar deep homologue). The first (which has already been done) is to insert (e.g.) an arthropod Pax-6/eyeless compliment into the mollusk camera-type eye development cascade, and see what type of eye develops; low and behold, if we substitute a Pax-6/eyeless from drosophila for its homologue in the octopode eye cascade, we get a normal octopode camera-type eye, *not* a hexapod compound eye (Gehring and Ikeo 1999).[4] The second manipulation, which has not yet been performed (and is perhaps more empirically challenging), is to replace the downstream, non-homologous generators of the octopode camera-type eye with those of the arthropod compound eye, while leaving the mollusk Pax-6 in tact—if an endogenous arthropod eye develops (a somewhat disturbing thought!), it is clear that the macroscopic arrangements of the eye (T) are causally determined by their downstream generators (P), which screen-off Pax-6 and other upstream homologues (D).

[4] The misexpression of the Pax-6 transcription factor has been shown to lead to the formation of differentiated ectopic eyes in both vertebrates and invertebrates. This data may seem to represent a counterexample to the claim that Pax-6 is casually insufficient for the formation of the macroscopic eye. However, this objection is neutralized by the contingent fact that the abnormal expression of Pax-6 simply triggers a cascade of downstream developmental events which are directly responsible for the substance and structure of the ectopic eye.

Theoretically, P (still largely a black box) will consist of a homogeneous (or maximally specific) reference class within which no statistically significant division can be made with regard to the production of T. This is in contrast to the heterogeneous class formed by taking the downstream generators in conjunction with deep homologues like Pax-6, as the partitioned probability of T given D is not equivalent to (and indeed a far cry from) the partitioned probability of T given P (which is equivalent to the non-partitioned probability of T given both D and P). While the precise reference class of P is currently unknown, there is good reason to believe that it exists and that it does not include Pax-6. Instead, the homoplastic work appears to be done subsequently to and independently of the master control sequence. It follows that insofar as their downstream generators are not homologous, vertebrate/mollusk/arthropod eyes represent convergent (not parallel) evolution.

In sum, just as T (a phenotypic component) may be said to generally screen-off P and D (genotypic components) with respect to reproductive success, proximal developmental mechanisms screen-off upstream homologues with respect to the production of T. Understanding parallelism in this way takes us from Simpson's operationally recalcitrant notion of a "community of common ancestry" to a bountiful research program in evolutionary development. Furthermore, it successfully deflects charges by

some that the difference between convergence and parallelism reflects human convention rather than important causal differences underlying different types of homoplastic regularity.

SCM's only mention of the parallelism/convergence distinction in his singular work on the philosophical importance of homoplasy is in a brief footnote, in which he states: "Now is the time to avoid that old chestnut of whether it is convergent evolution as against parallel evolution" (2003, 435 n.1). SCM reasons without argument that the difference is "obviously" one of degree rather than kind (Ibid). A central goal of chapter, however, is to convince the reader that the difference between convergent and parallel evolution is no small philosophical chestnut. SCM's failure to recognize the implications of this distinction for his hard adaptationist view of life is perhaps the most severe weakness in his appeal to the literature on homoplasy in order to bolster the project of a universal biology. In particular, conflating parallelism with convergence leads SCM to conclude that design space is more constrained than it actually is.

4. The Philosophical Implications of Homoplasy

Having established a meaningful conceptual distinction between convergence and parallelism, we can now address two key questions: First, how frequently do parallelisms occur, and second, how important are they

in understanding the history of life? Both of these questions in turn depend on our ability to identify instances in which two lineages converge (so to speak) on a morphology that is directly produced by a shared developmental apparatus. Because the proximate mechanisms of development, or what Carroll (2005, 110) calls genetic "dark matter," have long remained a black box, biologists have been compelled to infer parallelisms from general propinquity of descent. However, recent advances in evolutionary developmental biology have yielded some of the first clear-cut instances of macroscopic parallel evolution. One nice example is the parallel evolution of elongated or shortened pelvic spines in stickleback fish (Shapiro et al. 2004), an adaptive feat that has been accomplished independently numerous times by parallel changes in hindlimb development—specifically, with respect to a switch that effects the Pitx1 gene responsible for pelvic fin development (Carroll 2005). As the last Ice Age receded, populations of stickleback fish were isolated in glacial lakes, rapidly (<10,000 years) evolving parallel reductions and elongations in spines. They independently (and iteratively) assumed two basic ecotypes in response to common selective pressures: A benthic, short-spined form and a pelagic long-spined form. The former configuration reduces the chances of the fish being snagged by predatory dragonfly larvae, while the latter increases the diameter of the fish so as to exceed the gape of many

open-water predators. Other instances of parallel evolution have been documented in the spot patterning of *Drosophila*, presumably by changes in a switch responsible for the expression of a pigment-producing protein (Ibid). More recently, it has been shown that extant avians have retained the ability to develop archosaurian (crocodilian) teeth, a trait absent in birds since the end-Cretaceous (and lost independently several times in non-avian theropods) (Harris et al. 2006). This dormant developmental program is thought to be controlled by a signaling center at the oral/aboral boundary that controls the expression of teeth.

It is quite possible that regulatory changes in homologous downstream sub-circuits are causally implicated in many homoplastic events. If so, this would suggest that internal constraint is not limited to restricting the so-called 'evolvability' of lineages by reducing the overall isotropic variation on which natural selection can act, thereby rendering inaccessible large regions of morphospace; in addition, it may have a *positive* influence on evolvability by establishing preferred internal channels that make good solutions (including Good Tricks) more accessible to selection (see e.g. Gould 2002). Understanding the mechanical and/or developmental correlates of a primary adaptation may in turn allow for a degree of macroevolutionary predictability, at least within specified developmental parameters.

For example, Hunter and Jernvall (1995) have shown that the 'hypocone,' the extra cusp characteristic of the quadritubercular (rather than triangular) molars of most therian mammals, evolved "convergently" (read: in parallel) within this order of mammals on more than twenty occasions and is associated with marked diversification. It is almost certain that this highly iterative and presumably useful adaptation, which greatly increases occlusal area for crushing herbivorous material, owes its repeated evolution to a common developmental pathway underlying the generation of tubercles in the therian dental arcade. But it would be unreasonable to argue that the mammalian hypocone, because it evolved numerous times independently in one order of mammals, is an inevitable evolutionary outcome insensitive to developmental constraints. For even if parallelisms are, as Dennett suggests, nature's way of rewinding the tape of life, they only represent iterative outcomes of a very small rewind in a grand history of life replete with possible frozen contingencies.

Such philosophically deceptive regularities may underlie many of the paradigmatic examples of 'convergent' evolution that SCM invokes, which are better thought of as parallelisms (see Sterelny's (2005) review of Conway Morris (2003) for a similar criticism). For instance, SCM appeals to numerous examples of 'striking convergence' *within taxonomic classes* in response to common selective pressures. Such convergences, better

interpreted as parallelisms, include (*inter alia*): The 'convergent' suite of saber-toothed morphology between placental and marsupial felids (1998, 202-204; 2003, 130-132), 'convergences' in fossorial lifestyles within specific orders of mammals (132, 140), 'convergence' with respect to 'pike morphology' in several genera of freshwater fish (133), 'convergence' in raptorial forelimbs of the mantids and neuropterans (129) (two orders within the class hexapoda), 'convergences' within orders of birds in plume coloring, wing shape, and hummingbird morphology (138), and 'convergence' in stem morphology of the New World cactus and the African spurge (134), to name a few.

As to this last example, SCM states (Ibid) that "cacti and spurges are only distantly related, and their common convergence is because of the rigors of living in an arid environment." But this is only a partial truth, as a significant (if not dominant) causal determinant of the above recognized similarity is a common developmental architecture and (quite possibly) shared proximate generative mechanics. Similarly, the parallel elongation of canines and modification of supporting skeletal-musculature between marsupial and placental saber-toothed cats are likely underwritten by changes in conserved regulatory homologues at the relevant phylogenetic depths. Even if such instances represent genuine cases of convergence under my more restrictive definition of parallelism, less-than-proximate

developmental constraints—if frozen—still bias the set of potential adaptive solutions.

Pictorial renditions of these parallelisms are on their face quite remarkable, and might appear to reflect the inexorable power of natural selection to steer form toward certain pre-ordained evolutionary 'endpoints' or functional 'islands.' But they are not all that surprising or impressive, given a more thorough appreciation of developmental pathway conservation. As noted by Wake (1991, 555) in the context of the evolution of the attenuate body form in salamanders, the fact that related taxa have independently adapted to similar environments by evolving essentially identical ecomorphologies is suggestive not of formal-functional invariance but rather significant design limitations due to entrenched developmental pathways. SCM may be correct that natural selection will (*ceteris paribus*) tend to find the optimal solution to a given design problem, but it will be forced to do so within the strictures of contingently defined developmental parameters.

In sum, parallelisms with causal developmental homologues at shallower phylogenetic depths are not going to give SCM the kind of generalizability that he needs to support his universal biology, which requires a set of laws regarding the evolution of biological form that is *stable across higher taxonomic counterfactuals*. One cannot point to

examples of parallel evolution as evidence of strong, non-contingent constraints on adaptational design space, or to suggest the inevitability of certain evolutionary outcomes given a grand rewind of the history of life. That said, for the comparative physiologist, instances of marsupial-placental mammal homoplasy represent a "treasure of information distinguishing between crucial and incidental features *of mammals* that have taken up different habits and habitats" (Vogel 1996, 301, emphasis added). Non-accidental generalizations regarding the evolution of mammalian form can obtain even if one monumental contingency (such as a massive bolide impact at the K-T boundary) triggered the mammalian adaptive radiation in the Paleocene. This is because our macroevolutionary extrapolations could specify particular developmental constraints (which function as a limited set of *ceteris paribus* qualifications), thereby presupposing certain major transitions in the evolution of life or the origins of major (or minor) body plans. For an example of finer grain, given the peculiar haplo-diploid genetic system of Hymenoptera, we may one day be able to specify the ecological conditions that will tend to give rise to sexually non-egalitarian eusociality. The precision of macrobiological generalizations will tend to be inversely correlated with the taxonomic abstraction at which they are made—for instance, intra-ordinal extrapolations will fare better than those made within phylum and class,

due to the increased significance of shared homology which reduces the amount of isotropic variation that can serve as fodder for natural selection, and sets the parameters for relative ecological optimality.[5]

In order to rescue the philosophical import of many of his proffered examples of convergence properly reinterpreted as instances of parallel evolution, SCM could offer the following argument: The contention that conserved internal channels underlie many homoplastic events need not entail capitulation to a historicist, contingency-infested view of life, so long as each entrenched developmental component of any given homoplasy arose not by contingent but predictable, functional processes. Accordingly, so the argument goes, we arrive at the current inhomogenous distribution of morphospace occupation through the optimizing hand of selection working synergistically with the 'ratcheting' effect of internal developmental constraints. So long as the generators directly responsible for a given parallelism are themselves the product of globally optimal processes operating at deeper phylogenetic layers, parallelisms can form the basis of meaningful macroevolutionary generalizations.

---

[5] SCM does draw upon some examples of genuine convergence that could be the subject of robust macrobiological generalizations, such as those pertaining to sensory modalities (2003 Ch. 7); unfortunately, he offers no principled method for comparing the philosophical or inductive significance of different types of homoplasy. Additionally, he fails to show that any of the evolutionary endpoints that he infers from the distribution of homoplasy are associated with either diversification or persistence (or some other measure of evolutionary success), which (in my view) undercuts the notion that such outcomes represent stable islands of form amidst a roiling sea of stochasticity.

There are good reasons for rejecting this argument, however. While there are instances in which natural selection has been shown to act as an optimizing agent, simple engineering analyses demonstrate that selection acts more often as a tinkerer, tweaking pre-existing developmental schemes to meet transient local adaptive challenges. In fact, it is the ubiquitous sub-optimality of organismic design that is often promoted (by Darwin, no less) as among the best evidence for evolution by natural selection, a blind mechanistic process that works with what is at hand and cannot plan ahead. As Vermeij states (1994 223), organisms tend to embody "ad hoc and often rather clumsy solutions to functional demands, solutions that bear a deep stamp of history and ancestry" (citations omitted). If we take the concept of entrenchment seriously, it is difficult to believe that selection has produced globally optimal developmental machinery, particularly given the decaying nature of the selective environment.

Some of SCM's assertions tiptoe around the evolution of organismic form per se into the more Platonic realm of disembodied functional kinds. He argues that there are certain "biological properties" or "functional attractors" toward which form tends to gravitate, as evidenced by the pervasiveness of homoplasy. There are deep problems with this view, however, and precisely how deep will depend on how we cash-in the term

biological property. If we define it broadly to include such mechanisms as 'predator evasion,' 'metabolism,' 'thermo-regulation,' or 'propulsion,' such properties may indeed be universal but they would be of no use whatsoever in predicting the evolution of *form* if there is an unmanageably large disjunction of configurations that can realize the same function. The biological property of 'a behavior that increases one's chances of reproduction' would be about as uninformative. Similarly, although SCM draws upon adaptations relating to Batesian and Mullerian mimicry, these concepts in and of themselves offer nothing by way of the evolution of form because they are contingent on the sensory capabilities of predators, rather than any factors intrinsic to mimicry (Vogel 1996). Convergence between phyla or other distant taxonomic groups on certain broad functional solutions such as 'mastication' will not admit of interesting predictions regarding the evolution of form absent some specification of a particularized developmental framework that natural selection can tweak in furtherance of the common ecological task. It would be hasty (and indeed unimaginative) to conclude that vertebrate teeth, arthropod mandibles, or mollusk beaks are inevitable solutions to such a broad ecological design problem as 'chewing.'

Furthermore, homoplasy—SCM's data set of choice—is indicative of constraints on form, but it does not necessarily imply limitations on the

universe of function. Given the complexities of local, Earth-based ecology, it is easy to envision a near-infinite set of functional attractors in the hyper-heterogeneous *cosmic ecology*—one so large that science can never even make a dent in it, let alone formulate laws of form on its basis. The lack of specificity inherent in SCM's notion of 'biological property,' and its attendant conceptual and operational difficulties, prevents him from developing these ideas into a constructive research program in counterfactual biology.

SCM states (2003, 307) that "whenever the known edge of the evolutionary envelope is reached...then it will be explored independently several times." However, he does not offer any concrete guidelines as to how in any particular case we are to know when the putative "edge" is reached, nor does his unsystematic work on homoplasy take us any closer to realizing this goal. How many cases of convergence are to suffice, and how are we to determine the degree of similarity necessary to recognize them as philosophically meaningful homoplastic events? How do parallelisms fit into this philosophical project and how should they be delineated? Even more importantly, why are certain endpoints at the end? These questions confront any appeal to homoplasy as evidence for a non-contingently constrained adaptational design space. Whether the history of life turns out to be robust and repeatable, or quirky and contingent, the

nature of macroevolution is sure to have profound implications for some of the grandest questions in philosophy and science.

***

Figure 1

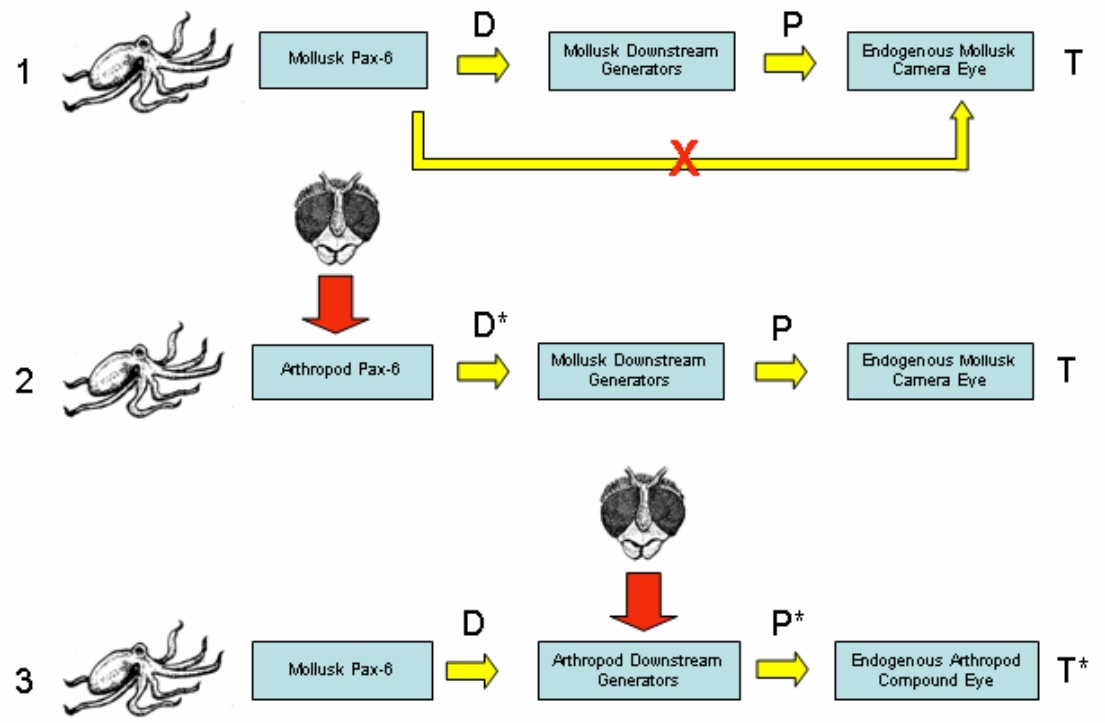

In the above diagram, three different scenarios are presented: In the first, a normal cephalopod Pax-6 compliment (distal cause 'D') triggers normal cephalopod downstream generators (proximate cause 'P') which produce a normal cephalopod camera-type eye ('T') that is homoplastic in relation to vertebrate camera eyes and (less so) arthropod compound eyes. It is clear that if either D or P is non-functional, T will not be produced. Although the literature sometimes speaks of Pax-6 as being both necessary and sufficient for eye morphogenesis, in actuality when expressed it simply regulates downstream generators which do the substantive work, such as crystallin genes which form the lens.

In order to show that P screens off D with respect to T, we do two manipulations, reflected in scenarios (2) and (3). In the first, we insert an arthropod Pax-6 compliment ('D*') into the mollusk camera-type eye development cascade, finding that instead of getting a hexapod compound eye, we get a remarkably normal cephalopod camera-type eye (demonstrating massive conservation of the upstream homologue). In the second, we leave the original cephalopod Pax-6 intact, but substitute arthropod downstream generators for their mollusk counterparts in the cephalopod developmental cascade. Presumably, the result would be an endogenous arthropod compound eye ('T*'), rather than T. Therefore, P may be said to screen off D with respect to the production of T.

Chapter 2

This Gouldian View of Life: Contingency and Convergence in Macroevolution

1. Introduction

For the last two decades, biologists and philosophers have debated whether the shape of complex life represents the fluky culmination of an eminently unrepeatable series of contingencies, or whether there is a greater necessity to life's grand parade of forms. This controversy was ignited by the publication of Stephen Jay Gould's *Wonderful Life: The Burgess Shale and the Nature of History* (1989), in which the late paleontologist invites us to contemplate the following evolutionary thought experiment: Rewind the 'tape of life' to when the first animals evolved (some 500 million years ago), and consider how its story would again unfurl. Gould believed that replaying life's tape would result in a radically different macroevolutionary outcome—a morphological menagerie bearing little resemblance to complex life as we know it.[1] Not only would no

[1] Gould invokes this thought experiment in an attempt to answer the following question: Is the clumpy distribution of organismic form in an otherwise vast and uncharted 'morphospace' the result of predictable optimizing processes, or is it the contingent upshot of quirky, unpredictable events which took place early in the history of life? (2002, 347) Although records of macroevolutionary transitions are plentiful, rarely if ever do we see a smooth gradation of forms bridging the body plans of the higher taxa.

humans, mammals, or vertebrates evolve, but neither would any creatures even remotely approximating them (1989, 291). For Gould, "almost every interesting event of life's history" falls within the realm of historical contingency (290). I will refer to this view of life as the *radical contingency thesis* ("RCT").

On the other end of the contingency spectrum is the view that re-running the tape of life would produce strongly similar (if not identical) macroevolutionary outcomes—beasts not unrecognizably different from those that have graced the Earth. I will call this view the *robust repeatability thesis* ("RRT"). It is based on the premise that the history of complex life can (for the most part) be attributed to the comparative selective advantage of evolutionary survivors over their sub-optimal, extinct brethren. On this view, successful animal phyla are not the fortunate winners of a macroevolutionary lottery, but a superior set of forms carved out of the vast morphospace of biologically possible but functionally suboptimal alternatives. The RRT does not claim that replaying life's tape will produce mammals *per se* at a particular space and

---

Presumably, such phyla-level transitions have been confined to the base of the Cambrian, leaving little record of their existence. Gould attributed this lumpy patterning of form to internal developmental constraints, rather than the optimality of current design under the unfettered operation of natural selection (2002, 1053). He viewed the set of actualized forms as a "subset of workable, but basically fortuitous, survivals among a much larger set that could have functioned just as well, but either never arose, or lost their opportunities, by historical happenstance" (1160-1161).

time, but it does argue that over immense spans of geological time, the evolutionary crank will tend to churn out eerily similar animal forms (Conway Morris 2003/1998; Dennett 1995, 307). Among the strongest evidence for the RRT is the independent origination of similar biological forms, an evolutionary phenomenon known as *convergence*.

The contingency debate has been encumbered by several key conceptual shortcomings. Perhaps most notable is the need to flesh out the notions of "contingency" and "convergence" in ways that are applicable to macroevolution. My argument shall proceed in three steps. First, I will propose a coherent, unified interpretation of macroevolutionary contingency. Second, I will show that in shoehorning all homoplasy into the category of convergence, many authors have inadvertently conflated what are empirically and philosophically distinct modes of iterative evolution. To remedy this, I offer a principled basis by which to distinguish parallel from convergent evolution, and I defend the distinction against recent challenges from developmental biology, in part by recourse to philosophical work on the ontological prioritization of biological causes. Finally, in light of the above reformulations, I show how the evidence frequently held to negate the RCT not only fails to do so, but in fact militates *in favor* of the Gouldian view of life.

John Beatty (2006) has recently argued that a particular brand of convergent evolution, namely 'iterated ecomorphology,' contradicts the Gouldian view of life, while it supports the view that natural selection carves out counterfactually stable patterns in macroevolution. I will use Beatty's appraisal of Gouldian contingency as an anchor (and not a punching bag) for my critique. Although the ensuing discussion engages problems particular to the philosophy of biology, the hope is that its conceptual machinery will prove useful elsewhere in the historical sciences.

2. The Radical Contingency Thesis

Gould believed that the strongest evidence for radical contingency comes from the fossil record of the earliest animals—in particular, the taxonomically intractable marine fauna which inhabited the Cambrian seas some 500+ million years ago (1989, 290).[2] This date marks the paleontological event known as the Cambrian Explosion, which refers to the geologically 'abrupt' origin of nearly all the major 'body plans' (phyla) in the animal kingdom, in addition to many fantastical designs that perished in the end-Cambrian extinction event (~488 mya). Among this

[2] Although Gould never abandoned the idea that there are important philosophical lessons to be drawn from the Cambrian fauna, it has become increasingly clear over the years that a number of these 'bizarre' taxa are less phylogenetically problematic than they initially appeared, many turning out to be stem groups of otherwise familiar phyla (Budd 2001).

motley menagerie, there is one taxon which, according to Gould, is of particular significance for its insignificance: This is *Pikaia*, a relatively primitive and understated creature in terms of its ecology and anatomical complexity, but one which by most accounts is the probable ancestor to all modern vertebrates. The moral Gould draws is this: Had conditions in the end-Cambrian been just a little bit different, then *Pikaia* (and the other fortuitous ancestors of living animals) would not have survived, and the subsequent morphological landscape would have assumed a markedly different shape (322-323). For Gould, the tale of the Cambrian fauna does not simply recount "a unique and peculiar episode of possibilities gone wild"—it betrays a profound truth about the nature of complex life itself (317).

Gould's argument for the deep contingency of animal life rests on two basic assumptions. The first is that the patterns of survivorship generated by the end-Cambrian extinctions are unrelated to the comparative fitnesses of the lineages. Consequently, the evolutionary coronation of some lineages and the extinction of others were governed by an effectively random rather than a robust and repeatable series of events.[3] The second premise is that once these early chapters in the Book of Life were written, they significantly constrained the form and content of

[3] By "effectively random," I simply mean for reasons unrelated to adaptation, fitness, or natural selection. I make no assumptions here about the underlying metaphysics.

subsequent chapters—much as the reader's initial decisions in an interactive 'choose-your-own-adventure' book disproportionately shape the possibility space of the journey. In other words, the fortunate survivors became *frozen* accidents due to 'contingent forces' that are (at least superficially) similar to those responsible for the universality of the genetic code (Crick 1968). What are these forces?

On the standard historicist account of macroevolution, morphological entrenchment does not reflect a stable and all-things-considered optimal evolutionary solution.[4] Instead, it is attributed to 'phylogenetic inertia' resulting from a peculiar property of complex developmental systems. Specifically, once the genetic networks responsible for the overarching morphological parameters of an organism are laid down, they become highly impervious to perturbation. This is due to their being causally bound-up with myriad interacting genes and functional pathways, making them highly resistant to drift and directional selection. The result is a developmental network that cannot be modified atomistically, with the more elaborate mutations necessary to modify it

[4] In explaining the origin of a biological feature, historicists tend to privilege "passive inheritance" or "phylogenetic inertia" over the utility or optimality of current function (Gould 2002, 1052). A third macroevolutionary weltenschauung is 'structuralism,' which attributes the ubiquity of independent similarity not to the supremacy of selection, but to non-functional (indeed, non-biological) laws of complexity (e.g. Kauffman 1995).

being vastly improbable, due to the random nature of genetic variation.[5] Diversification is thus limited (so the argument goes) to the confines of the specified 'body plans' of a few surviving lineages of the early mass extinctions (2002, p. 1159). In sum, by combining fitness-independent survivorship at higher taxonomic levels with the 'inertia' of developmental constraint, you arrive squarely at the RCT.[6]

On its face, the RCT dovetails nicely with the received view in the philosophy of science that biology operates in what amounts to a nomological vacuum (Rosenberg 2001, 1985; Beatty 1995). The lawlessness of biology, and the exception-riddling of even its best generalizations, is thought to stem from several peculiar features of the subject matter: (1) the multiple realizability of function, (2) the supervenience of fitness on stochastic properties of the environment, and (3) the unrelenting arms race

---

[5] Although there is evidence for a surprising degree of developmental robustness due to canalization, buffering, and dominance mechanisms which can accommodate genetic perturbation (Wagner and Schwenk 2000), *expressed* variation in radically conserved regulatory networks will rarely if ever be sustained. Embryonic development is highly unstable against perturbations in transcription factors which affect cascades controlling cellular differentiation (Erwin 2006; Thattai and van Oudenaarden 2001). This is because, as gene disruption studies indicate, the phenotypic effects of genetic disturbances are not linear or modular; in many cases damage to the phenotype includes not only the structures that are directly implicated by the mutation, but 'collateral' traits as well.

[6] The word 'inertia' here is somewhat misleading, since unlike Newtonian physical systems which tend to stay the same unless acted upon by an external force, biological systems have a tendency to change (i.e. drift) unless acted upon by natural selection (Brandon 2006). Given that change (not stasis) is the null expectation in biology, powerful selection pressures may be necessary in order to preserve the phylogenetic status quo.

of strategic co-evolution. Given these characteristics of evolution, it is unlikely that any morphological regularity will exhibit the nomic necessity characteristic of the physical laws. This has important implications for the contingency debate, which concerns the counterfactual resilience of not merely what organisms and their phenotypic traits *do* (i.e. their selected effects or causal role functions), but rather *how* they do it. That is to say, the debate is one about *form*, not merely function. While biological lawlessness does not rule out (and the arguments in this paper remain agnostic to) the existence of meaningful *functional kinds*, the above features of biological evolution would seem to foreclose the possibility of natural *morphological kinds*.

3. The Challenge from Convergent Evolution

Increasingly, however, biologists and philosophers are pointing to convergent evolution—or the independent origination of similar biological forms—as evidence against the Gouldian view of life. Because all life on earth shares a single common ancestor, and because we do not currently have recourse to an extraterrestrial data set, one might reasonably question whether the issue is ripe for anything beyond fumbling speculation. Yet many authors regard convergent evolution as tantamount to independent experimental replication in the history of life. Gould

acknowledged as much, but he maintained that convergence was a relatively unimportant phenomenon in the macroevolutionary scheme of things (2002 , 1068). To the contrary, Dan Dennett (1995) and Simon Conway Morris (2003) have both touted convergent evolution as strong evidence not only for adaptation, but also for *hard adaptationism* (*sensu* Amundson, 1994), or the view that nearly all scientifically interesting features of life can be explained by natural selection. As these authors have been critiqued elsewhere (e.g. Powell 2007), this paper will focus on John Beatty's recent evaluation of the RCT. As Beatty is not himself a hard adaptationist—and because he is in fact generally sympathetic to Gouldian themes—his critique shows that this reading of convergent evolution is not limited to the more extreme advocates of hard adaptationism.

As a paleobiologist, Gould spent more time building empirical support for the RCT than he did fleshing out its conceptual underpinnings. Hence Beatty's recent paper titled "Replaying Life's Tape" (2006), which examines both the conceptual and empirical dimensions of the controversy, is a welcomed contribution to the debate. Nevertheless, in what follows I will show that several key conceptual shortcomings lead Beatty to misinterpret the nature of macroevolutionary contingency and (consequently) the evidentiary implications of convergence. My critique

shall proceed in two steps. First, I will evaluate (and ultimately reject) Beatty's pluralistic conception of macroevolutionary contingency, and I will defend an alternative, unified account of the RCT. Second, in light of this re-formulation, I will consider the implications of convergent evolution for the Gouldian view of life.

4. The Nature of Macroevolutionary Contingency

Beatty convincingly argues that Gould equivocates between two compatible but importantly different conceptions of contingency. I will discuss each of them in turn and show why they fail, taken either individually or collectively, to capture the essence of macroevolutionary contingency. My aim here is not to quarrel with Beatty over his interpretation of Gould, nor is it to defend the RCT against its detractor theories; rather, it is to come up with a unified notion of contingency that gels with Gould's larger theoretical framework—regardless of whether one subscribes to that framework or not, and notwithstanding any rhetorical ambiguities that may have invited a more pluralistic interpretation.

The first conception of contingency that Beatty attributes to Gould is "contingency as causal dependence," which implies that a series of prior events in a chain of causation are each necessary with respect to the production of an outcome ('O'), such that if any of these events had not

occurred or occurred in a different way, O would not have occurred or would have occurred in a different way. At best, this definition of contingency is over-inclusive, as it fails to rule out nomically expectable outcomes: If all events along a causal chain are highly likely to repeat (say, due to constraints of the laws of physics), then O will be virtually certain to repeat, given a replay of the system. At worst it is trivial, insofar as it entails the metaphysical platitude that *some* change in initial conditions will tend to produce *some* change in outcome. As it stands, contingency as causal dependence is unable to distinguish between events as different as an asteroid-induced mass extinction and the eight ball falling predictably into the corner pocket. No one denies that if an object the size of Mars crashed into the Earth during the early Cambrian, the subsequent history of life would have been markedly different. But this is not the crux of the contingency debate.

The second type of contingency that Beatty ascribes to Gould is "contingency as unpredictability," which entails that identical initial conditions do not suffice to produce the same outcome. This definition seems to accord with the 'rewind the tape' thought experiment, whereby we go back in time to the early Cambrian and let life march forward once again, only to find that it does so to a very different macroevolutionary tune. On its face, this notion of contingency would seem to commit Gould

to metaphysical indeterminism, since it requires that the *same initial conditions* produce *disparate outcomes*—a physical impossibility if determinism holds for biological systems. And yet, the inference from contingency as unpredictability to indeterminism is one which Beatty expressly disavows (2006, 345), and rightfully so given that Gould explicitly divorced randomness from contingency (1989, 283). On the other hand, if determinism obtains then rewinding the tape is a trivial exercise, for it will always play out in *precisely the same manner*.

To sidestep these metaphysical difficulties, Beatty must explicitly exclude from the initial conditions certain stochastic features of genetics and development—namely, the generation of variation, ordering of mutations, and other trappings of "chance" which serve to underwrite evolutionary unpredictability. Beatty is puzzled by Gould's decision not to include stochastic processes in his concept of contingency, given that "Gould acknowledged [these phenomena] as sources of historical contingency" (345). But simply because drift is a casual *source* of contingency does not entail that it is a *type* of contingency. Moreover, I do not think (nor do I believe that Gould thought) it entirely implausible that irreducibly probabilistic (e.g. quantum mechanical) processes could influence mutational trajectory and, in turn, macroevolution. Nonetheless, in order for Beatty to absolve Gould of any such metaphysical

commitments, in referring to disparate outcomes from "the same initial conditions," he must either be excluding a large and important set of boundary conditions, or else referring to an *epistemically equivalent* set of the same, wherein negligible genetic and environmental differences are responsible for the disparity in outcomes.

Beatty offers a slightly different formulation of contingency as unpredictability which he also ties to Gould, one which "denies that evolution by natural selection is sufficient to guarantee the same evolutionary outcome, even given initially indistinguishable ancestral lineages and indistinguishable environments, and even excluding stochastic processes like genetic drift" (339). Here it looks as if Beatty is again referring to the idea (which he attributes to Gould) that stochastic processes (like mutational ordering) will be a necessary part of any evolutionary explanation. One problem with this definition of contingency is that it conflates the robustness of macroevolutionary pattern with the nature of its underlying mechanism(s). It is a common mistake to assume that the predominance of selection is antithetical to radical contingency. But even if a macroevolutionary trend can be explained by natural selection alone, it may be eminently unrepeatable if it is generated by many complexly configured selection vectors that are distributed randomly with respect to one another (Millstein 2000). For instance, if each species

in a lineage went extinct due to unique selection pressures, resulting in the extinction of the entire clade, the clade-level pattern could not be attributed to a single adaptive story, even if it is overdetermined by selection. Thus, there is no a priori reason to suppose (as Beatty assumes Gould does), that the prevalence of selection would imply or even be positively correlated with the robustness of macroevolutionary pattern.[7] In other words, the contingency debate does not boil down to the respective significances of selection and drift in macroevolution. The question is not whether natural selection is at the helm of macroevolution, but whether it knows (metaphorically) where it is headed.[8]

---

[7] In this context, 'robust repeatability' relates to a particular *outcome*, not to the nature of the *processes* which produced it. This is not to say that underlying mechanisms are causally irrelevant to robustness, but merely that they are not constitutive of it.

[8] As Beatty correctly suggests, any notion of contingency should rule out a robust equilibrium explanation of macroevolution, in which disparate starting points lead inexorably to a single attractor (as exemplified by the bowl-and-marble system). Perhaps a better physical illustration of the sort of robustness at stake in the contingency debate is something like a large, dense cloud of particulate in the void of space: A vast number of possible particle distributions will produce a singular outcome—namely, a mass rounded by its own gravity (e.g. a planet or star). This is not to say that a high energy collision could not prevent this from occurring; but again, the contingency debate is less concerned with scenarios in which ceteris is far from paribus. If we rewound the cosmic tape to the early moments of the universe and let it play out once again (holding the laws of physics constant), it is virtually certain that stars, planets, black holes, pulsars, and other familiar celestial entities would evolve—since they supervene on an enormous range of spatio-temporal configurations. While the particular distribution of matter in the universe is highly contingent on early boundary conditions and chancy quantum events, the properties of material inhomogeneity and their predictable celestial consequences are counterfactually robust. Likewise, the contingency debate is not about the origins of specific taxa at

In sum, Beatty's account of macroevolutionary contingency does not get at the philosophical heart of Gould's hypothesis, and this in turn causes him to misread its empirical application. I will now offer an alternative, unified conception of contingency which I take to be closer to the metaphysical core of the Gouldian view of life: I will refer to it as *radical contingency*. The origins of radical contingency can be found in some of Gould's earliest writings on the subject, where he describes the quintessential case of contingency as one in which "small and apparently insignificant changes…lead to cascades of accumulating difference," yielding entirely different evolutionary outcomes (1989, 290). But consistent with Beatty's ambiguity thesis, Gould can also be found associating contingency with "an unpredictable sequence of antecedent states, where *any major change* in *any step* of the sequence would have altered the final result" (emphasis added). I suggest that in order for the contingency concept to do the work that Gould intended it to do, it must entail changes of particular magnitudes and at particular stages in the history of life.

Broadly construed, radical contingency is the notion that arbitrarily small differences in input variables produce disproportionately great

---

particular locations in the history of life (Dennett 1995, 307)—events that are preceded by millions of complexly configured antecedent states. The question, rather, is whether there are any biological forms which, like their cosmological counterparts, exhibit a wide range of counterfactual invariance.

disparities in outcomes. Thus, *outcome O in system S is* ***radically contingent*** *iff a marginal change in some initial condition* $I^1 \ldots I^n$ *of S would tend to result in Outcome O*, where O* is radically disparate from O*. The key point is this: Marginal disparities in initial conditions tend to lead not only to different—but *radically* different—evolutionary outcomes. Gould hints at this interpretation of contingency when he suggests by way of analogy that if we were to rewind the tape of the American Civil War, "with just a few small and judicious changes (plus their cascade of consequences), a different outcome, *including the opposite resolution*, might have occurred with equal relentlessness" (1989, 283, emphasis added).[9]

This formulation of Gouldian contingency incorporates the valuable elements of Beatty's pluralistic reading, while eschewing many of its difficulties. First, radical contingency does entail a causal dependency, but it is a *particular sort* of causal dependency—namely, an outcome's sensitive dependence on marginal changes in initial conditions. In other

[9] This notion of radical contingency is not distinctively biological, for it applies equally to weather systems and stock markets as much as it does to organisms and taxa; but in this respect it is no different from Beatty's notion of contingency as causal dependence. Radical macroevolutionary contingency, on the other hand, entails the above physical dynamics in relation to a particular set of evolutionary outcomes—namely, those at or above the species level. Even if macroevolution is chaotic, however, this does not preclude the existence of certain morphological 'attractors' whose origin and stability is probable across many rolls of the evolutionary dice. Whether convergent evolution is evidence for the existence of such attractors is the question to be taken up in subsequent sections.

words, it implies a *chaotic* causal dependence. Second, we can understand "contingency as unpredictability" in light of this super-sensitivity to boundary conditions. Chaotic dynamics can magnify arbitrarily small differences in evolutionary environment, resulting in disparate outcomes from initial conditions that are otherwise metaphysically or epistemically identical. In addition, this analysis of contingency helps to frame the contrast class. Contingency is usually couched in opposition to repeatability; but recall that if determinism is true, then the tape of life would be eminently repeatable, and hence a trivial thought experiment. It is therefore not repeatability per se—but rather *robust repeatability*—that characterizes the antipodal view, where robustness relates to the stability of an outcome over a wide range of initial conditions.

It follows that in order to determine whether a particular evolutionary outcome is contingent or robust (or somewhere in between), we must (1) identify the relevant class of initial conditions, (2) delineate outcome similarity space, and (3) specify perturbation magnitudes. For the sake of brevity, I will focus on (1) and (3), and simply assume (arguendo) that (2) has been met.[10] How should we circumscribe the relevant class of

---

[10] I will nonetheless offer some preliminary thoughts on outcome specification. One can imagine numerous methods for carving up the relevant regions of morphological outcome space. In the context of the history of life, we might ask why primates, why mammals, or why vertebrates evolved, rather than something else? But what is that *something else*? Such contrast classes are conceptually

initial conditions? For starters, we can subtract all matters of fact that have no statistical bearing on the outcome (such as the position of Saturn). This still leaves us with a mind-boggling array of initial conditions whose manipulation would affect the probability of the outcome. Additionally, we must decide what kind of change in initial conditions should constitute a *marginal* one? Is a three-degree change in temperature to be considered 'marginal'? What about a magnitude seven rather than a magnitude five earthquake? Although these questions are not easy to answer, much of these complications can be avoided through reasonable stipulations. From the vantage point of the RRT, mundane perturbations in the history of life (such as 'run-of-the-mill' changes in temperature or plate-tectonic activity) should not radically and permanently alter the morphological landscape of life on Earth. Such events may be anything but 'mundane' insofar as a single clade is concerned (what is marginal for an order may be devastating for an individual species), but they should be considered

---

elusive, not so much because of the subjectivity of the delineation, but rather due to the fact that the space of morphological possibility is both vast and (for the most part) uncharted. Even more problematic is the list of threshold morphological conditions necessary to establish something like 'mammal-hood.' What does it mean to say, as Conway Morris (2003) does, that 'mammal-ness' is a robust biological property? For this to be so, not only must each adaptation in the constellation of traits for mammal-ness be robust on its own, but it must be necessarily co-extensional with all of the others. Consider, for instance, the amniotic egg, feathers and bipedalism: These traits are an integral part of 'birdness,' and yet there is no logical (or biological) reason to expect their co-occurrence. Likewise, the camera-type eye, convergent in vertebrates, cephalopod mollusks, and arachnid arthropods, is an adaptation that seems to stand on its own, although it need not co-occur with an endoskeleton, mantle-secreted shell, and jointed exoskeleton, respectively.

marginal for the purposes of assessing radical contingency, which is a thesis about the overarching shape of life.

5. The Philosophical Implications of Homoplasy

Having fleshed out the concept of macroevolutionary contingency, we are now in a better position to evaluate Beatty's claim that convergent evolution contravenes the RCT. Here, Beatty joins other philosophers and biologists who have appealed to the independent origination of similar biological forms (for which the technical term is *homoplasy*) as evidence that natural selection has shaped robust and repeatable patterns in the history of life (Vermeij 2006; Conway Morris 2003; Dennett 1995). Beatty is rightfully skeptical about the testability of Gould's (or the antipodal) thesis, given that a single instance of contingency or robustness cannot falsify a biological world view. He asks "what sorts of studies, short of a complete tally of evolutionary episodes, will give us more than anecdotal insight into the overall importance of historical contingency?" (2006, 362 fn. 2). What Beatty fails to realize, however, is that even an exhaustive inventory of "convergent" events in the history of life will not be dispositive. In lumping different sources of homoplasy together, Beatty overlooks an important distinction: Namely, that between parallel and

convergent evolution ("the P-C distinction"), each a type of homoplasy but with different underlying causes and ensuing philosophical implications.

When systematists are faced with a similar trait in two distinct lineages, they ask whether the similarity is homologous (i.e. due to common descent), or whether it is homoplastic (i.e. derived independently). Typically, this question is answered by reconstructing the phylogenetic relationship of the two lineages (in the form of a 'cladogram') to determine if their common ancestor exhibited the trait in question. If so, then the trait is deemed a homology; if not, then it is considered a homoplasy—end of story. However, I submit that for the purposes of the contingency debate, there is one more step that needs to be taken: Namely, to ascertain whether the homoplasy is an instance of *convergent* or *parallel* evolution. This distinction will prove relevant to evaluating the relationship between homoplasy and the RCT.

Homoplasy between closely related lineages is often referred to as 'parallelism,' while that between more distant groups has generally been designated as convergence (for a review of the terminology, see Arendt and Reznick 2007, 28). Since all known living things are related to some degree or another, one might be inclined to think that the convergence-parallelism distinction tracks an irreducibly spectral phenomenon—namely, relatedness—that can only be partitioned arbitrarily. Thus, many authors

(Beatty included) have tended to either ignore the distinction between parallel and convergent evolution (lumping all homoplasy under one category or the other), or to conclude that that the difference is ultimately one of degree rather than kind (Abouheif 2008; Arendt and Reznick 2007; Diogo 2005; Conway Morris 2003, 435 fn. 10).

I submit that there is in fact a non-arbitrary, scientifically operational, and philosophically important distinction between these two types of homoplasy. In a nutshell, my contention is that *a homoplasy is a parallelism just in case a developmental homology is the* ***proximate*** *cause of the phenotypic similarity*. It is true that some authors have associated parallelism in two lineages with a common developmental substrate that has been retained since their divergence from a common ancestor, but that is unexpressed in their *most recent* common ancestor (Hall 2007; Meyer 1999). But all such definitions fail to take into account causal asymmetries in trait development, and thus encourage the perception that parallelism and convergence spill over into one another. To the contrary, the above definition is objective in that it picks out a natural, causal dividing-line between superficially similar but fundamentally distinguishable evolutionary events. Furthermore, it is amenable to empirical intervention to determine whether a given homoplasy is of one type or the other. In particular, a 'screening-off' test can be used to ascertain whether a given

homologue is the proximate cause, rather than simply *a* cause, of a given homoplasy.[11]

The P-C distinction is important for the contingency debate insofar as it suggests that ostensibly 'independent' macroevolutionary replications are not so independent after all.[12] Although the connection between convergence and robustness has not been made explicit, RRT proponents appear to view homoplasy as corroborative of their thesis insofar as it implies that a highly *dissimilar* set of forms arrived at a highly *similar* set of morphological outcomes (via the optimizing forces of natural selection, which are treated as constant). At the most basic level, homoplasy entails that two lineages L1 and L2, exhibiting similar form T, share a common ancestor in which that form was absent. Presumably, this absence is due to the lack of certain developmental generators responsible for the relevant

---

[11] The idea is simple: Proximal genetic cause P screens-off more distal cause D (e.g. a shared master control gene) of homoplastic trait T where the probability of T given P and D, is the same as the probability of T given P, and different from the probability of T given D. Just as T (a phenotypic component) may be said to generally screen-off P and D (genotypic components) with respect to reproductive success, proximal developmental mechanisms screen-off upstream homologues with respect to the production of T. See Powell (2007) for a more elaborate discussion of the screening-off relationship in the homoplasy context.

[12] In making use of proximate developmental homology, this operational definition of parallelism is a significant improvement on Gould's loose engineering metaphor in which he compares Pharaonic bricks to Corinthian columns—the former being present in all existing structures, the latter shaping the peculiar organization of a particular architectural tradition (2002, 1138).

phenotype. The working (and deeply flawed) assumption is that the visible disparity in initial morphological conditions maps onto or is commensurate with a similar disparity in underlying developmental conditions—something that is often not the case (see below).

6. Defending the Parallelism-Convergence Distinction

Recently, developmental biologists have cast considerable doubt on whether the P-C distinction can be maintained, given that it relies on a linear, oversimplified model of evolutionary development which rarely obtains in nature. In this section, I show that the P-C distinction (as above conceived) is indeed vulnerable to such objections, but that it can nonetheless be salvaged by recourse to philosophical work on the ontological prioritization of biological causes.

The most rigorous and persuasive rebuttal of the P-C distinction is due to Arendt and Reznick (2007), who view the concept of parallelism as an unhelpful relic of a previous age when evolutionary biologists were not privy to the underlying developmental causes of phenotypic variation and, as a consequence, failed to appreciate the complexity of the genotype-phenotype map (31). Their contention that the P-C distinction should be abandoned rests on two major premises. The first (and less compelling of the two) relates to the fact that *closely* related lineages can evolve the

same phenotype via *different* developmental mechanisms, while *distantly* related lineages can derive a similar feature via *identical* genetic substrates which have been retained in latent form since their separation from a common ancestor. But even if we accept the authors' stronger molecular-evolutionary assertion that "there is no predictable association between taxonomic affinity and similarity of the genetic basis for [a given homoplasy]" (p. 30), this poses no real difficulties for the above definition of parallelism. The first scenario describes a case of genuine convergence (i.e. homoplasy produced by wholly different developmental substrate), while the second a classic case of parallelism (i.e. homoplasy produced directly by a latent developmental homologue). We might think of the former as an example of 'shallow convergence' and the latter an instance of 'deep parallelism,' given the respective propinquities of the comparison groups. But no harm is done to the distinction in either case.[13]

[13] This first rationale for rejecting the P-C distinction falls short for another reason. In light of the frequent 'decoupling' of genotype and phenotype, it will not suffice to identify a developmental disparity underlying a given homoplasy and brand the latter as convergence on that basis, as Arendt and Reznick have done. New genes can appropriate previously unrelated developmental pathways without any resultant break in morphological continuity. In such cases, trait homology is preserved despite a complete turnover in developmental mechanics (Roth 2001, 94). Even if a similar morphology has evolved in two lineages 'independently,' and even if it arises from wholly distinct developmental mechanisms, it can nonetheless be classified as parallelism so long as it exhibits the requisite *developmental continuity*. Thus, in order to determine on which side of the line a given homoplasy falls, we must look not only to current molecular function but also to genealogy.

The authors' second reason for jettisoning the P-C distinction poses a more serious threat to my thesis. Previously, I argued that a homoplasy is a parallelism if and only if a developmental homology is the *proximate cause* of the phenotypic similarity, and I proposed a 'screening-off' test to ascertain whether an identified homologue is the proximate cause (rather than simply *a* cause) of a given homoplasy. The trouble with this analysis, however, is that it only seems to work in the context of a Markov-like causal sequence leading from the genotype to the phenotype. Yet, as Arendt and Reznick show (2007, 30-31), few morphological traits will be generated in such simplistic topological fashion, given epistasis and the non-linear interdependencies of gene networks. Topological nonlinearity may represent the rule rather than the exception for biological systems (Wagner 1999). Some authors have even questioned whether we can speak meaningfully of causality (or at least *regular* causality) in the context of qualitatively nonlinear biological systems (Ibid, 94). While the screening-off test can still be used in such cases to show that proximate developmental generators screen-off more distal genetic homologs with respect to the production of a homoplastic trait, the parallelism claim hinges on the additional assumption that the lineages sharing a homoplasy also share the proximate developmental mechanisms (i.e. regular causes) which produced it. And if Arendt and Reznick are right, it is unlikely that

any independent morphological similarity will be produced by *identical* developmental mechanisms, since such traits tend to emerge from interacting gene networks involving an array of distinct genetic causes. The upshot is that homoplasy will rarely meet the clean-cut definition of parallelism.[14]

Before attempting a rebuttal, we can further strengthen Arendt and Reznick's argument by drawing out some of the counterintuitive implications which flow from the generic nature of the proximate cause criterion. If we accept that parallelism is a spectral concept due to varying degrees of overlap in the proximate developmental cause, then we are forced to regard as parallelism the scenario in which two homoplasy-bearing lineages share *nothing more* than a proximate 'accessory' protein. On this view, a homologue could be the sole basis of parallelism even if it does no *substantive morphological work*. Arendt and Reznick (2007, 30) conclude that rendering a similarity judgment about the underlying

[14] As Arendt and Reznick report, even simple cases of parallelism will often involve complex networks of genes with differential pleiotropic effects. For example, although the same amino acid polymorphism (Mc1r) has been associated with pigmentation gain in various mammals (from beach mice to wooly mammoths), its function affects or is affected by non-homologous developmental components which are (in each case) equally necessary for the development of the trait. A similar lesson can be drawn from Pitx1, which is associated with pelvic alterations not only in sticklebacks (discussed in section 6), but also in more distant vertebrate clades—playing a key role, for example, in the fin-to-limb-to-fin transition in marine mammals (Shapiro, Bell and Kingsley 2006) and possibly even reptiles (Caldwell 2002). Does this suggest that these ostensible convergences are parallelisms after all? Not on Arendt and Reznicks' view, since any homologous genetic cause will generally be integrated with non-homologous ones that are jointly necessary for trait production.

developmental mechanics of a homoplasy is akin to "divining between shades of gray rather than discerning black from white." For simplicity's sake, they argue, we should refer to all homoplasy as "convergence." One might go even further and argue that since there is no hard and fast P-C distinction, there is no non-arbitrary basis on which to claim that one instance of homoplasy is any more (or less) compelling than another, insofar as the contingency debate is concerned. Both of these conclusions are erroneous.

For a homoplasy to constitute a parallelism of *any magnitude*, there must be at least *partial* homology with respect to the proximate developmental cause of a homoplastic trait; convergence, on the other hand, entails that there is *no relevant homology* in the same. But partial proximate developmental homology is merely a necessary condition for parallelism—it is not sufficient. What matters is not homology per se, but homology at the *relevant causal depth* ***and*** of the *relevant causal type*. The former refers to the proximate developmental cause that is identified by screening-off manipulations; the latter picks out a specific causal ontology that will require some unpacking.

Recall that the conserved pigmentation gene responsible for parallel adaptive coloration in mammals (Mc1r) is actually bound up with various non-homologous gene networks in the distant lineages in which it is

expressed (see fn. 15, above). The question we are confronted with is this: is our ability to identify a shared developmental cause of the pigmentation homoplasy undermined by non-homology in the complex developmental pathways which produce it? Phrased in more general terms, does the existence of nonlinear causal networks in biological development imply that there is no non-arbitrary basis on which to privilege some genetic causes rather than others?

This is essentially the question taken up in a recent paper by Ken Waters (2006), in which he extends Woodward's (2003) counterfactual theory of causation into the realm of developmental biology. Ontologically speaking, not all causes are created equal, as Woodward has shown. We have good metaphysical reason not only to distinguish between causes and non-causes, but also to pick out the *actual* difference-makers from the vast set of *potential* difference makers in explaining the variation in outcome. This enables us to say that Mary's striking the match is an ontologically distinct cause of it lighting in one case and it not lighting in another, or of it being unlit at time T and lit at T+1. Other causes, such as the presence of oxygen or phosphorus, do not vary across scenarios in which there is an *actual empirical difference* (though of course they could). On this view, DNA sequence is *the actual* cause of RNA structure in a bacterium, even though RNA polymerase and other accessory proteins are *necessary* causes

as well. This is because actual differences in DNA explain the actual variation in RNA sequence, while the accessory proteins do not vary. But in cases where DNA *and* accessory proteins vary, both are actual causes of RNA structure.

Unfortunately, this takes us no closer to saving the P-C distinction. Woodward's philosophical apparatus allows for the ontological privileging of *statistically relevant* and *actually differing* conditions in explaining variation across a population of outcomes. In the present case, the actual difference to be explained is the character state of the homoplasy shared by two lineages versus that of their most recent common ancestor. The question we must ask is this: Is there a common developmental cause that is *actually and exclusively* responsible for the difference in character state? If Arendt and Reznick are right, then the answer will generally be 'no'—at best, there will be only partial homology with respect to the proximate developmental cause of a homoplasy. Hence, the utility of the P-C distinction is completely undermined.

Thankfully, Waters (2006) gives us a way out. Building on Woodward's general theory, he makes a persuasive case for causal asymmetry in biology—not only with respect to potential and actual difference makers, but also between *specific* and *non-specific* actual causes. Specific causes are those processes which, if subjected to a battery of

interventions, will tend to change the outcome in numerous and detailed ways. Nonspecific causes, on the other hand, merely determine *whether* or *when* an outcome will occur; they have no influence on precisely *how* it will do so. To spell out precisely what this means, recall the above discussion of RNA synthesis. On Woodward's account, there is no basis on which to assign causal priority to DNA over and above accessory proteins with respect to the construction of RNA, so long as we assume that both are *actual* (rather than merely potential) difference makers in relation to RNA variation. On Waters' reading, however, DNA is *the specific* actual difference maker, since alterations in DNA engender particular changes in RNA sequence, whereas changes in accessory proteins are limited to halting the synthesis process entirely or merely altering the rate at which it takes place. But DNA is not the only specific actual difference maker with respect to RNA molecule variation. As Waters recognizes, RNA splicing agents, which remove *particular* segments of RNA and fuse the remaining segments together, are also causally specific. Thus, we can privilege certain biological causes over others given the circumstances of particular cases, but we cannot privilege a priori one *class* of biological entities (e.g. DNA) over another (e.g. ribonucleoproteins).

Waters' thesis was formulated in the context of RNA transcription, where rate is only temporally relevant to the outcome. One problem with

extending this line of reasoning to morphogenesis is that alterations in the rate and timing of development—a phenomenon called 'heterochrony'—can have profound morphological and evolutionary consequences (for an overview, see Gould 1977). As such, ontogenetic factors controlling for the rate and timing of trait development are not limited in causal scope to Waters' "whether/when" category, since they influence morphological parameters and other substantive factors that fall within the "how" dimension of the outcome. At the same time, variability in the amino acid sequence of proteins can have little or no impact on morphology (consider 'isozymes,' or structurally different enzymes that catalyze the same chemical reaction).[15] There is ample reason, therefore, to be skeptical of any attempt to rank biological phenomena from the get-go.

Waters' philosophical machinery has important implications for the present discussion. To see why, let us return to the question of parallel mammalian pigmentation. The P-C skeptic contends that because hundreds of interacting, multiply deployed genes underwrite even such simple adaptations as pigmentation gain, there will be little similarity in the developmental pathways that lead to parallel coloration. Even if an important sequence (such as Mc1r) is shared, the proximate developmental pathways of any given homoplasy will be largely non-homologous. Are we nevertheless justified in claiming that a particular segment of DNA (or a

[15] For this and the previous point, I am indebted to V.L. Roth.

distinct transcription factor) is ontologically privileged over other genes, gene products or regulatory elements which are equally necessary for trait production? If so, this would enable us to single out those developmental factors that make an actual specific difference to the trait in question, and this would in turn rebut the charge that parallelism is a sloppy and hopelessly subjective category.

In order to identify the specific actual difference-maker of a homoplasy for the purposes of assessing parallelism, we need to look for a homologous sequence of structural DNA containing instructions that are ultimately translated into the synthesis of the specific proteins that, with the aid of regulatory components, determine the relevant morphological parameters.[16] So long as such a homologue is present in both lineages and is a substantive determinant of the gross morphology of shared trait T, it does not matter (for the purposes of parallelism) whether T also relies on non-homologous accessory proteins or functionally unrelated structural DNA for its production. For the same reasons, homology in regulatory (i.e. cis-acting) regions of the genome, such as promoters, enhancers, silencers,

[16] Because the loss of a function (such as pigmentation) can be effected by the alteration of many different genes and processes along the tortuous route from genotype to phenotype, homology in such cases will often fail to pass the threshold of specificity needed to establish a parallelism on my account of the P-C distinction. The implication is that many instances of independent trait loss represent convergent rather than parallel evolution. Given the myriad ways in which the synthesis of a trait can be obstructed, it seems reasonable to conclude that parallel gains are less probable than convergent losses.

and other factors affecting gene expression, will often (but not always) be an insufficient basis on which to ground the specific internal difference-maker. Unlike nucleic acid sequences, transcription factors do not code for proteins or RNA polypeptides. That said, as noted earlier, changes in developmental rate and timing can be of profound morphological consequence. It would be a serious mistake, therefore, to say that regulatory elements cannot in principle do any substantive morphological work.[17] So long as the developmental trait is *homologous* (inherited from a common ancestor) and *causally specific*, it can form the basis of a parallelism. It is an open empirical question whether homologues like Mc1r exhibit the kind of specificity needed to rise to the level of a specific actual cause. The point I am making, however, is a theoretical one: It is

---

[17] For instance, consider the ontogenetic basis of 'Polyphenism', which refers to the generation of alternative phenotypes in response to differences in their developmental environment. The polyphenic threshold, or the conditions under which the shift between alternative phenotypes occurs, can evolve via natural selection by changes in hormonal regulation—in particular, through the canalization or sensitization of a plastic phenotype. Suzuki and Nijhout (2006) showed that 'sensitizing' mutations in hormonal regulatory pathways can reduce hormonal titers, thereby decreasing the polyphenic threshold and allowing for the expression of otherwise hidden variation under conditions of environmental stress. Once the hidden norm of reaction (due to the accumulation of phenotypically 'silent' mutations) has been exposed, mutations in modifiers can then alter the post-embryonic threshold, resulting in substantial morphological evolution. The revelation and selection of hidden variation via parallel modifications in hormonal regulatory pathways is a powerful modus operandi for parallel evolution. But rather than a simple on-and-off switch, a complex network of sensitizing mutations, coding sequences, and environmental fluctuations are jointly responsible for the actual differences in polyphenic traits.

not the *extent* but rather the *causal type* of developmental homology that counts.

Finally, were we compelled to discard the P-C distinction, this would not preclude us from adjudicating between different *tokens* of homoplasy with respect to their evidentiary implications. Greater degrees of proximate developmental homology indicate that the initial conditions are more similar than one might have surmised on the basis of morphology alone. By the same token, the absence of proximate homology can transform ostensibly minor homoplasies (such as those between closely related taxa) into much more impressive examples of convergence than they would otherwise seem to be. The thrust of my argument depends not so much on the P-C distinction per se, as it does on the implications of different sorts of homoplasy, whatever the accepted terminology.

7. Iterated Ecomorphology as Evidence *against* the RCT

Having wrestled (hopefully successfully) with the core concepts of the controversy, we turn now to consider the evidence. In putting the contingency debate to the test, Beatty reviews several evolutionary 'experiments' designed to investigate whether evolution is contingent on unique past events, or whether directional selection will lead disparate populations to converge on a common adaptive solution irrespective of

their histories. He devotes a fair amount of discussion to the ingenious laboratory experiments of Travisano et al. (1995); for the sake of brevity, however, and because Gould's contingency thesis was formulated in connection with animal form, I will confine my critique to examples that implicate *macro-morphological* evolution. This is not to say that the evolution of microbial metabolism is a trivial feat; but at bottom the contingency debate concerns the broadest brush strokes on the canvas of organismic form—not the capacity to digest maltose.

The first putative counter-example to the RCT that Beatty cites is the iterated adaptive radiation of the Canadian threespine stickleback fish (2006, 338 fn. 2, citing Schluter 1994). As the last Ice Age receded, populations of stickleback fish became isolated in glacial lakes. With remarkable rapidity (in less than 10,000 years), they independently and iteratively segregated into two ecomorphs in response to common selective pressures. The first is a benthic (bottom-feeding) short-spined form, and the second a pelagic (open-water) long-spined form. The former configuration reduces the chances of the fish being snagged by predatory dragonfly larvae, while the latter increases the diameter of the fish so as to exceed the gape of open-water predators. On its face, this looks like a clear-cut example of robust repeatability in macroevolution.

A closer look supports a different interpretation. The stickleback radiations are among the first clear-cut and thoroughly documented instances of parallel macromorphological evolution. Evolutionary changes in pelvic armor of these closely related populations have been accomplished independently numerous times by parallel regulatory changes in a single conserved Mendelian factor, either by way of recurrent mutation or persistent polymorphism (Foster and Baker 2004). The expression of this single latent homolog of major effect is directly responsible for the parallel ecomorphology of globally distributed populations of sticklebacks (Shapiro et al. 2004). While such parallelisms may be indicative of repeatability *per se*, they fail to provide empirical support for *robust* repeatability. For if a simple matter of gene regulation is the only difference-maker in terms of initial conditions, and if all relevant structural genes are conserved, then the similarity in outcome is not all that surprising, given the similarity in initial conditions. This makes parallelism in general—and the evidence that Beatty cites in particular—look more like the trivial version of contingency as mere causal dependence, rather than a decisive counterexample to the RCT.

The second data set apparently inconsistent with Gouldian contingency is the Caribbean anole lizard radiations, which along with the sticklebacks represent some of the best-documented examples of iterated

ecomorphological evolution. As many as six distinct 'ecotypes' have evolved repeatedly and independently on isolated islands in the Greater Antilles (Losos et al. 1998).[18] According to Beatty, this series of independent experimental replications speaks in favor of robust repeatability and (by logical implication) against the RCT. Yet like the stickleback radiations, the anole lizard phenomenon is highly susceptible to a parallel evolution explanation, as even the researchers themselves concede.[19] Beatty relates that Gould was not much impressed by iterated ecomorphogenesis, maintaining that the RCT concerns taxonomically deeper evolutionary counterfactuals (Beatty 2006, fn. 16). But the contention that homoplasy at shallower phylogenetic depths cannot speak for or against the RCT is, on my view, both inadequate and incorrect—but not for the reasons that Beatty and others might think. It is *inadequate* because it relies on an unwarranted focus on phylogenetic depth per se, rather than on the nature

[18] The ecotypes vary in features including limb-length, skull dimensions, and other traits relating to predator escape and foraging ability. For instance, species occupying open habitats tend to have long legs for increased sprinting ability, while those inhabiting branches have shorter legs which increase their maneuverability in this specialized adaptive zone. Despite their considerable morphological disparities, all within-island populations of lizards are phylogenetically closer to one another than to any inter-island population.

[19] To date little is known about the developmental biology of the anole lizards, and specific genes associated with changes in hind-limb development, skull morphology, skin pigmentation, and other traits that comprise the various ecomorphologies have yet to be identified. Nevertheless, researchers in the Losos lab believe that key developmental homologs exist and will ultimately form a crucial part of any synthetic explanation of anole radiations (Sanger et al. 2007). Thus it is safe to wager that like the sticklebacks, the anole radiations represent parallel (rather than convergent) evolution.

of genetic variation and its causal relation to macro-morphological evolution. It is *incorrect*, because there are in fact ways in which homoplasy confined to the lower taxa could detract from the plausibility of the RCT. For instance, the stickleback and anole scenarios would cut against the RCT if it turned out that the various ecomorphs were generated from disparate developmental substrates, rather than the "flickering on and off" of latent regulatory homologues of major phenotypic effect in response to similar selective regimes (Abouheif 2008, 3). But truth be told, if we are to truly shake the Gouldian view of life, we would need to see genuine convergence *across* (not just *within*) the higher taxa. With a few notable exceptions (like the image-forming eye), this is simply not the case.

Despite their high degree of developmental affinity, the various stickleback and anole clades did indeed begin their evolutionary journey from different (albeit not radically different) starting points. As Beatty argues, if the RCT were correct, then we would expect a *dramatic divergence* rather than a *narrow convergence* between these closely related lineages. It would seem, then, that neither example supports either of the competing hypotheses. In the next and final section, however, I will argue that the above parallelisms do in fact support Gould's thesis, which is more nuanced than many authors (including Beatty) have recognized.

8. Iterated Ecomorphology as Evidence *in favor of* the RCT

Beatty assumes that if the RCT is correct, then independent adaptive radiations should lead to disparate evolutionary outcomes, even if the starting conditions are similar (2006, 305 fn.2; *accord* Losos et al. 1998, 2115). On its face, this interpretation seems to mesh well with my interpretation of radical contingency, according to which small differences in initial conditions lead to large discrepancies in outcome. To make the case that iterated ecomorphology affirmatively *supports* the Gouldian view of life, I will have to delve somewhat deeper into Gould's theory.

Gould invoked the 'tape of life' thought experiment in the context of explaining the inhomogenous distribution of organismic form in a theoretically vast morphospace (see fn. 2 above). He asked whether the clustering of variation around a coherent, stable set of body plans reflects the ecological excellence of those designs vis-à-vis their extinct competitors, or rather the unique and unrepeatable signature of history. Throughout his career, Gould vigorously defended the latter, arguing that patterns at the grandest scale of animal evolution can be explained in large part by internal developmental constraints on the evolution of form. Once the Cambrian extinctions had culled the initial crop of body plans,

large regions of evolutionary possibility were rendered permanently off-limits, leaving gaping holes in morphospace that were never re-occupied.

But there is also a positive side to the story, one which Beatty and other critics have overlooked. On Gould's view, animal evolution has been "positively abetted (as much as negatively constrained) by homologous developmental rules acting as potentiators for more rapid and effective selection" (2002, 84). Although developmental networks are generally resistant to perturbation, when they *are* disturbed they tend to shift in a few preferred directions. Because only few mutations *of phenotypic significance* can be had without catastrophically undermining developmental integration, evolutionary trajectory will tend to bend towards the region of morphospace linked to those mutations. This allows internal constraints to work *synergistically* with directional selection, providing a reliable conduit for fitness-enhancing change. As Gould states, "homologous developmental pathways can also be employed [] as active facilitators of homoplastic adaptations that might otherwise be very difficult, if not impossible, to construct in such strikingly similar form from such different starting points across such immense phyletic gaps" (2002, 1122-1123). Once established, this bias in development allowed for the iterative activation (or cooptation) of the same genes of major effect in response to analogous ecological design problems (Abouheif 2008, 4). For

Gould, these realities of evolutionary development (which had just begun to emerge at the time of his later writings) are consistent with his allegedly heterodox conclusion that macroevolution is driven by "top-down channeling from full ancestral complements, rather than [the] bottom-up accretion along effectively unconstrained pathways of local adaptation" (2002, 84). Rather than a monolithic apology for radical contingency, Gould's theory entails *local pockets of predictability embedded in and casually dependent on a larger framework of radical contingency*.

Gould explicitly anointed parallelism as the sine qua non of this 'positive' dimension of internal constraint (2002, 1122-1123). Notwithstanding his occasionally superlative rhetoric, Gould did not view parallel evolution as an *alternative* to Darwinian gradualism, but instead as a theoretical bridge between micro and macroevolution (*accord*, Abouheif 2008). Rather than slam-dunk evidence for the power of natural selection, Gould attributed parallelism to the 'congealing' of ancient developmental machinery. He did not deny that natural selection will tend to find the "Good Tricks" in design space, as Dennett (1995, 308) puts it; but he maintained that the reason why there are so *few* good tricks, and why these are so *readily accessible* to selection, is due to the internal channeling of developmental constraint which aids and abets evolutionary reiteration (2002, 1178). Gould maintained that both the Cambrian

Explosion and the post-decimation diversification within phyla owe their existence to positive internal constraints (84). The fact that Gould's thesis issues *seemingly* contradictory predictions (i.e. repeatability and contingency) may be a sound reason for rejecting it; after all, a theory that predicts everything explains nothing. But the goal of this paper is not to vet the empirical status or logical coherence of Gould's view of life. It is simply to show that parallelisms of the sort cited by Beatty do not undermine, and in fact reinforce, the logical structure of Gould's evolutionary theory.

9. Conclusion

In conclusion, by selectively focusing on a few remarkable instances of parallelism, many authors appear to have missed the forest for the trees. The vast majority of clades that have undergone multiple independent radiations under similar ecological conditions *have not converged* on a morphologically similar set of outcomes. Homoplasy may be the closest thing to independent experimental *replication*, but if so, then the history of life is replete with independent experimental *non-replications*. For instance, although benthic and pelagic lake habitats are commonplace, I am not aware of any evidence that the stickleback 'solution' has been replicated in other isolated clades of freshwater fish. The stickleback and

anole phenomena are of particular scientific interest precisely because they are *rare*. This suggests that there is something peculiar to their phylogenetic history that makes their particular solution a good one. For all of these reasons, we should be loath to generalize from a few instances of parallelism to robust replicability in the history of life.

***

Chapter 3

The Evolutionary Biological Implications of Human Genetic Engineering

1. The Evolutionary Harm Argument against Human Genetic Engineering

In 2006, Apple Computer® launched an ad campaign touting the virtues of Macs while lampooning the common foibles of Microsoft® PCs. The first commercial in the series, entitled "Viruses," portrays the back and forth banter between a sneezing man who represents a PC that has been infected with a virus, and another who symbolizes a Mac computer that is immune to the PC's 'cold.' By highlighting the fact that Macs are less susceptible to virtual viruses, the commercial implies that they are somehow "better designed" than their PC counterparts. To the contrary, however, the increased vulnerability of Microsoft computers is due not to any particular design flaw, but rather to Microsoft's enormous success in the computing world. Comprising over 90% of the operating system market, Microsoft software presents a target-rich environment for would-be virtual assassins. So much so, in fact, that the Computer and Communications Industry Association recently warned that Microsoft's dominance has

created a silicon-based 'monoculture', one that could spell security disaster for economic sectors which rely heavily on the Microsoft platform.

The term 'monoculture' has become increasingly pejorative in recent years, particularly among the ranks of environmentalists, anthropologists, and other vehement critics of globalization. But it has more congenial roots in the context of agricultural practice, where it refers to the growing of a single cultivated crop (or 'cultivar') over a relatively large swath of land. Because of the high genetic relatedness of the cultivars in a monoculture, their planting, maintenance and harvesting can be standardized, increasing the efficiency of crop production and (consequently) reducing the cost of food. As it turns out, however, the benefits of monoculture come at a substantial price—namely, an increased risk of catastrophic crop failure. Genetic uniformity in agricultural practices increases the chance of crop loss for two chief reasons: first, high levels of relatedness increase the vulnerability of a cultivar population to large-scale epidemics, which can spread rapidly in genetically homogenous populations; and second, low levels of biological diversity can impair the flexibility of cultivar lineages to respond to changing environmental conditions, such as fluctuations in temperature, moisture level, or soil composition.

Perhaps the most famous illustration of the perils of monoculture is the Great Irish Potato Famine of the middle 19th century, which led to the

death of nearly ¼ of the Irish population. The proximate biological cause of the potato epidemic was a single-celled, host-specific infectious organism (*Phytophthora*) that has been linked to numerous plant pathologies, including (and especially) potato blight. But a deeper explanation of the tragedy makes use of *evolutionary* biological facts: namely, that in planting clones of the 'lumper' potato variety in vast numbers and over wide areas, farmers unwittingly reduced host species diversity (Bourke 1993). In so doing, they effectively rolled out the genetic red carpet for this voracious eukaryotic parasite.

A similar but more recent anecdote relates to the Californian winery debacle which occurred near the end of the 20th century, and from which an analogous precautionary moral can be drawn. Years before the catastrophe, agricultural experts at the University of California (Davis) had recommended that wine-makers in the Napa Valley region use a productive rootstock cultivar called AxR1. This cultivar was thought to be resistant to the insect pest *phylloxera*, which had single-handedly wiped out nearly all the vineyards of 19th century Europe (Campbell 2004). As it happened, however, while AxR1 did retain its original resistance, the aphid-like pest had evolved into a form that could thrive on the AxR1 monoculture. This oversight, in addition to a lack of appreciation for the

dangers of host crop uniformity, led to the replanting of two million acres of vines, resulting in a financial disaster to the tune of one billion dollars.

The moral of the monoculture story can be read in two different (though not mutually exclusive) ways: "know thy mortal ignorance," or "know thy evolution." Regardless of the chosen emphasis, the basic message is clear: it is dangerous to put all of your agricultural eggs into one genetic basket. Why should the same precautionary maxim not apply with equal force to the genetic modification of humans, a technology which (ostensibly) threatens to narrow the range of human genetic variation? Critics contend that given our unfortunate experiences with monoculture, the burden of persuasion should be on those who seek to demonstrate the safety of human genetic modification, rather than on those who merely purport to identify its risks. I disagree with this allocation of the rhetorical burden, but I believe that the arguments in this paper will rise to the challenge in any case.

In a certain sense, there is nothing new in the idea that reproductive technologies and social practices could combine to decrease human genetic diversity, either in the aggregate or in any subset. This might happen, for example, if it became increasingly common to choose a mate or to abort a pregnancy on the basis of information obtained through genetic screening. But these technologies and practices could not result in

anything even approaching a monoculture scenario, since they do not affect background rates of recombination and mutation, the two primary sources of genetic variation. However, the same may not be said for *robust* genetic technologies, such as gametal genetic engineering, which can alter the genome—and by implication the gene pool—to an extent and with a degree of efficiency that is unprecedented in the history of life on Earth.

Thus far, the ethical analysis of germ line genetic engineering technology ("GET") has focused primarily on its social, psychological, or aesthetic-moral implications (see e.g. President's Council 2004/2002, Kass 2002/1998, and Habermas 2001/2003, respectively). Rather than re-tread this well-worn territory, I will concentrate on a challenge to GET that is commonly advanced but which has received far less critical attention in the literature: namely, the argument that GET will reduce the range of existing human genetic variation ("HGV"), creating a biological monoculture that could not only increase human susceptibility to disease, but even hasten the extinction of our species. Insofar as this paper explores the *phylogenetic* implications of GET, it compliments a recent paper in which Powell and Buchanan (forthcoming) examine the *ontogenetic* ramifications of the same technology. Although both papers consider GET in an evolutionary biological context, Powell and Buchanan focuses on the

development of traits during the lifetime of an organism, whereas the present paper is concerned with the evolution of human populations.

As I see it, there are two major areas of evolutionary concern which, taken together, comprise what I will refer to as the 'evolutionary harm argument' ("EHA"). Both components of the EHA hinge on the premise that GET will substantially reduce HGV. The first contends that a progressively homogeneous human population will become increasingly susceptible to disease e.g. (Rifkin 1983); the second claims that a shrinking range of biological diversity will reduce the human species' flexibility in responding to novel adaptive challenges (Baylis and Robert 2004). In broad form, the EHA entails that the regulation or blanket prohibition of GET is necessary to protect the diversity of the human gene pool and, by implication, to prevent the aforementioned evolutionary harms.

I will show that once properly fleshed out, the EHA is unpersuasive, since it is premised on several key misconceptions about the nature of genetic variation and its relationship to phenotypic diversity, disease resistance, evolvability, and the mechanism of natural selection. In this paper, I argue that the widespread use of GET is unlikely to reduce HGV, and that even if it did, this would neither increase the human species' susceptibility to disease, nor prevent it from responding effectively to the shifting demands of selection. By rejecting GET in order to preserve the

health of humanity and its valued characteristics, we may be jettisoning the most powerful weapon in our adaptive arsenal for ensuring the long-term survival of our species (see Buchanan 2008a).

2. The Nature of Biological Variation

Thus far, the EHA has proven difficult to vet due to a lack of theoretical and empirical specificity. In order to cure this defect, we need to get a firm grip on the nature of biological variation. The presence of ample, heritable variation is a crucial premise in Darwin's 'one long argument' for descent with modification. When we speak broadly of 'human variation,' we are referring to all of the characteristics that make people different from one another, including traits that are culturally transmitted. *Biological* variation is a particular subset of human variation that refers to any and all genotypic and phenotypic diversity that is biologically transmitted. At the genomic level, measures of diversity include the number of alleles per locus or the overall proportion of genetic polymorphism; at the populational level, diversity is measured in terms of character trait variance; and finally, at higher taxonomic levels, diversity is indicated by species number, functional differentiation, or morphological disparity.[1]

[1] It is important to note that variation is not the same thing as *variance*, which refers to the distribution of variation around a mean. One population might have a large amount of variation tightly clustered around the mean, while another might have a smaller amount with a wider distribution in variation space. It

Darwinian evolution requires that heritable variation be the *cause* of a propensity for differential survival and reproduction. For the most part, natural selection acts directly on an organism's phenotype, and only indirectly on its genotype (Hull 2001; Brandon 1990). Because selection tends to operate at or above the organismic level, it only 'sees' the functional phenotype, and thus it is insensitive to the genetic substrate from which that function is realized. It stands to reason that HGV is important for adaptive purposes only insofar as it has, or will at some future time have, a tangible effect on the phenotype.

Because the EHA is typically couched in terms of genetic rather than phenotypic variables, the first thing we need to do is to consider the relationship between genotypic and phenotypic diversity. Philosophers tend to focus on HGV because they assume that phenotypic variation maps neatly onto genotypic variation. But in doing so, they succumb to the 'gene-for' fallacy, or the idea that each gene codes for a single trait and (conversely) that each trait arises from the operation of a single gene. The landscape of the genotype-phenotype map is actually far more complex, for several reasons.

The first is *phenotypic plasticity*. The phenotype is a product of the genotype and its interaction with the grab-bag category we refer to as the

---

could turn out that the range of existing variation, sometimes called *disparity*, is a more significant factor in disease resistance and evolutionary flexibility than the sheer volume of diversity itself.

'environment.' Because many phenotypic traits are highly plastic, they will develop disparately in dissimilar environments despite their underlying genetic identity. A single genotype can produce an array of phenotypes, each varying in accordance with the environmental context in which it unfolds (Via et al. 1995). The result is phenotypic diversity without a corresponding level of genotypic diversity. For example, consider the pupae of eusocial insects (such as ants, bees and wasps). These undifferentiated larvae kin, despite their high genetic similarity, can develop into members of the worker, soldier, or queen castes depending on the temperature, nutrition levels, and other environmental factors in which they are reared. The upshot is that high levels of phenotypic diversity can be maintained in a population without correspondingly high levels of genetic diversity.

The second is *multiple realizability*. Not only are we unable to infer much about genotypic diversity on the basis of phenotypic diversity alone, but the reverse also holds true. Many phenotypes are multiply realizable in that they supervene on a range of underlying genotypes. Natural selection will treat all variants equally so long as they have the same effect on the phenotype. Consequently, phenotypic uniformity can overlay substantial amounts of genetic diversity.

The third is *pleiotropy*. This one-to-many relationship, effectively the inverse of multiple realizability, describes the situation where a single

gene produces a wide range of functionally unrelated phenotypes. Pleiotropy is different from phenotypic plasticity in that the resultant trait diversity is due not to environmental heterogeneity, but rather to compound gene function. But like phenotypic plasticity, pleiotropy allows phenotypic diversity to supervene on genetic homogeneity.

The fourth is *nonlinearity*. Because of the complex causal dynamics of the genotype-phenotype map, changes in genetic sequence will rarely have a linear or proportionate effect on the phenotype. In some instances, small genetic perturbations can have enormous ontogenetic consequences. For instance, mutations that occur early in ontogeny (i.e. 'upstream' in the developmental cascade) can be amplified in the unfolding of the organism (Davison and Erwin 2006; Carroll 2005). In other cases, large genetic disturbances can be of minor phenotypic significance. Some functions are so well-buffered against developmental noise and genetic error that even large perturbations do not affect the resulting phenotype; in addition, large portions of the genome are non-functional, and thus can be modified without altering the phenotype.

Each of these phenomena is discussed in greater detail below. For now, what matters is that because of the non-symmetrical mapping of traits onto the genome, phenotypic diversity cannot be reliably inferred from genetic diversity, and vice versa. Failing to causally connect-up HGV

with phenotypic diversity, and the latter with natural selection, is one of the major oversights of the EHA. Another is that it lumps all types of genetic variation under a single generic heading. This conflation poses a problem for several reasons. First, nuclear DNA is only one type of genetic material that is transmitted into the next generation. The sub-cellular organelles, such as the mitochondria, possess their own genetic code as a relic of their free-living prokaryote days. It is unclear how this type of DNA would bear on any of the phenotypic traits that bioethicists care about.

But simply excluding the genes of organelles does not solve the conflation problem. This is because the nuclear genome itself is not a homogeneous reference class for the purposes of evolution by natural selection. The category of nuclear DNA can be further broken down into two different types of genetic diversity. The first is *neutral genetic variation*, which refers to genotypes that are orthogonal to or have no bearing on fitness; the second is *adaptive genetic variation*, which describes genes that are actively under selection (Kimura 1983). Given that this distinction is rarely acknowledged outside of the biological literature (Holderegger, Kamm and Gugerli 2006), it is not surprising that it is entirely absent from philosophical discussions of the evolutionary implications of GET.

In diploid organisms, or those containing two homologous copies of each chromosome, three different genotypes can occur at a given locus (e.g. aa, ab, bb). If the locus is non-adaptive (i.e. neutral), then it does not matter for the purposes of selection which of these genotypes is present, and the locus will accumulate mutations stochastically. If the locus is under selection, however, then it does matter which variant is present, and selection will eliminate the relatively less fit ones, thereby reducing genetic diversity at that locus. The fact that selection tends to *reduce* variation poses an ostensible paradox for Darwinian theory, since descent with modification requires a steady stream of variation to draw upon in response to changing environmental conditions. There is still much controversy as to the mechanisms that maintain genetic diversity in natural populations. Research over the last few decades, however, points to neutral variation as a critical ingredient in, and genetic drift as a central mechanism of, biological variation. This may sound counterintuitive, for while drift tends to increase variation between populations, it is generally thought to decrease variation within them by bringing certain variants to fixation (assuming the presence of absorbing boundaries). But in portions of the genome that have no effect on fitness, diversity can accumulate at a steady rate over time, thanks to mutation, drift, and other stochastic forces that go 'under the radar' of natural selection. These non-adaptive genetic

sequences have been given the (misleading) sobriquet 'junk DNA', and appear to constitute the vast bulk of protein variation (Nozawa, Kawahara and Nei 2007; Reich et al. 2002). When we choose any two people at random from the entire human population, we find that 99.9% of their DNA is identical. Of that 1/10 of 1% of remaining variation, a large proportion (~70%) is effectively neutral. To put it crudely, the majority of human genetic variation is junk.

In contrast to junk DNA, which has only captured researchers attention in the last few decades, adaptive genetic variation has been the focal point of evolutionary thought since its inception in 1859. In practice, however, adaptive genes are more difficult to identify than their neutral counterparts. This is because adaptive variation is inferred from patterns of complex traits, most of which are produced by nonlinear, epistatic interactions of gene networks. These complex developmental dynamics make it extremely difficult to infer levels of adaptive genetic variation from observed phenotypic diversity (Conner and Hartl 2004). Were adaptive and neutral variation correlated, this would provide a tractable means for measuring the former. But no such correlation has been revealed, and junk DNA cannot be used as a proxy for adaptive diversity.

Selection will tend to purge less fit variants from the gene pool, while neutral sequences will accumulate mutations steadily over

evolutionary time. In fact, it is the absence of expected variation that is the most reliable indicator that a gene or trait is under selection. It follows (somewhat counter-intuitively) that change, not stasis, is the null expectation in biology (Brandon and McShea 2008). Unlike Newtonian physical systems, which when at rest tend to stay at rest unless acted upon by an external force, biological systems have a tendency to change (i.e. drift) unless acted upon by natural selection (Brandon 2006). It follows that biological diversity should not be viewed as a goal to be achieved or a state to be actively maintained, but rather as an inherent disposition of replicating systems. GET may act to reduce genetic variation and thereby offset the propensity to drift, but in this respect it is no different than natural selection.

3. Will Genetic Engineering Technology Reduce Human Biological Diversity?

Having sketched out the landscape of biological variation, we are now in a position to consider the likely impact of GET on human genetic diversity. As noted in the previous section, one of the major shortcomings of the EHA is that it focuses on genetic variation *per se*, rather than partitioning this class into the causally differentiated categories of neutral and adaptive

variation. This conflation is more than a simple oversight—it amounts to a fundamental flaw in the EHA, for several reasons.

First, although the EHA touts the value of diversity, it is abundantly clear that not all biological variation is desirable. This may seem all too obvious, given that the very business of natural selection is to weed out unfavorable variants from the population. But the idea goes deeper than this. Beyond a certain age, humans will contribute little to the gene pool of the next generation, and thus (with some rare and controversial exceptions) natural selection will tend to ignore the post-reproductive period of life. Consequently, as the human organism ages, it invests less and less in the physiological repair mechanisms that would otherwise eliminate harmful genetic variation. Like a neglected house left to fall into disrepair, the body begins to accumulate genetic and ontogenetic variation, leading to disease and eventually death. Surely we do not desire the kind of genetic variation that leads to functional disintegration, such as that wrought by cancerous cell lines, neural degeneration, or recessive diseases. Thus, to make its case, the EHA must zero-in on the beneficial subset of variation, while excluding the diversity associated with conditions that we would treat as pathology.

Second, because the vast majority of HGV is neutral, and since biological systems will continue to accumulate variation in the absence of

selection, it is unlikely that GET (targeting phenotypes like eye color or attention span) will have a significant effect on the overall level of genomic diversity. Recall that in biology, diversity arises 'for free' in systems that are not under selection. For obvious reasons, GET will be geared towards engineering traits that make a difference to consumers of the technology. It will not waste time modifying unexpressed genetic sequences that have no palpable effect on the architecture or function of the organism. For this reason, GET will leave the lion's share of genetic diversity intact.

But even if we remove junk DNA from the equation and focus only adaptive variation, it is unlikely that GET would have a greater homogenizing effect than ordinary background selection. Although adaptive variation comprises a smaller fraction of the genome than junk DNA, at any given moment the number of genes that are under selection is vast. Even if we did manage to homogenize a subset of adaptive variation, the impact on overall functional diversity would be negligible. Those who think otherwise tend to overestimate the degree of genetic homogeneity that can be inferred from casually observed phenotypic traits. As studies in the biology of race have shown, the variation *within* putative racial groups is greater than the variation *between* them (Cavalli-Sforza 1994). Everyone in a society could look like either Ken or Barbie, and yet their underlying genetic diversity could rival that of any two randomly selected people on

earth. The set of traits that human beings tend to notice is but a tiny fraction of existing phenotypic variation.

Third, even if we assume that GET will lead to uniformity in a wide range of *phenotypes*, this need not entail a corresponding uniformity in their underlying *genotypes*. As we saw in the previous section, the same phenotype can be produced from disparate genetic substrates, given the many-to-one and one-to-many dynamics of the genotype-phenotype map. This is especially true for complex traits, the prime targets of GET, which rarely correlate with and only with a specific subset of the genome (Nijhout 2003). The implication is two-fold: first, the targeting of a particular phenotype by GET need not result in the homogeneity of its underlying genotypic generators; and second, the targeting of a particular genotype need not increase the uniformity of its protein-product (given *epistasis*, or the interaction between regulatory networks in relation to their effect on the phenotype). For example, we can increase phenotypic variance in the domestic dog population, producing an astounding array of forms from the Chihuahua to the Newfoundland, while at the same time decreasing total genetic diversity.

Fourth, even if GET did produce temporary pockets of genetic uniformity, whether they would be maintained is highly contingent on human population structure and the extent of gene flow between natural

populations. Revolutions in transport and information technology have led to unprecedented levels of global exchange, not only in relation to goods and services, but with respect to genes as well. With the exception of the occasional un-contacted Amazon tribe discovered accidentally by loggers, there are few behaviorally or geographically isolated human populations. As a result, any homogenization due to GET will likely be dampened and ultimately swamped out by invading variants. This scenario is particularly compelling, given that access to and usage of GET will be far from uniform, allowing localized pockets of homogeneity to be easily re-absorbed into the genetic mainstream.

Finally, even if GET did bring certain genotypes to fixation, causing the extinction of competing alleles and hence a reduction in overall genetic diversity, this would not be irreversible. In the wild, extinction represents a true absorbing boundary, particularly in the case of complex functional pathways whose iterated independent origin is extremely improbable. By contrast, human-initiated gene banks (akin to the Global Seed Vault which recently debuted in Norway) can be maintained, and from which genes can be retrieved, long after their extinction in the wild. Extinct genotypes can be 'resurrected' (as it were) in order to introduce favorable variants into the population or control for levels of genetic diversity. In conjunction with other reproductive technologies, such as nuclear transfer cloning, GET

could be used to facilitate the rapid re-deployment of genes (Buchanan 2008b).

The factors I have been discussing thus far are all biological. But whether GET is likely to increase or decrease human biological variation, and the extent to which it will do so, turns not only on biological facts, but also on the psychological, social, and political framework in which the technology is used. Broadly speaking, the impact of GET will depend on the nature of the genetic technology at issue, its demographic penetrance, the extent of individual/cultural convergence in use, and the existence of regulatory regimes that constrain its proliferation or function.

Let us begin by distinguishing cloning, or the crude duplication of an existing genome, from GET, which involves the precise manipulation of existing genes. In terms of its affect on levels of HGV, the pervasive cloning of a small number of individuals lies on one extreme end of the homogeneity spectrum. But even in this most extreme and unlikely scenario, it is perfectly possible to limit cloning to the *functional* components of the genome, while allowing for background diversity in neutral DNA. In this way, even the widespread cloning of a small subset of individuals could preserve a substantial proportion of existing HGV. It could turn out, of course, that the evolutionary value of non-functional DNA is negligible (a proposition that I contest in the final section); but the

point is that one need not clone the entire genotype in order to reproduce the same phenotype. On the other hand, if cloning technology was both accessible to and utilized by a wide range of persons, then the reductions in HGV would be far less severe. The higher the penetrance of cloning technology, the less impact it would have on human biological diversity. For instance, if every living human cloned him/herself only once at time T, then the resulting genetic pool would be no more or less diverse at time T+1, and presumably no more or less susceptible to risks associated with homogeneity than the existing human population.

Nevertheless, most authors assume that access to sophisticated reproductive technology will, at least initially, be limited to the wealthy, thus skewing the gene pool in favor of the upper echelons of society. This is the crux of the skeptic concern—namely, the mass production of a small number of genetic types. But it fails to take into account the strong negative correlation between income level and expected reproductive fitness. Despite their superior resources, richer people tend to have fewer children than those of the less privileged classes. This forces the EHA to overcome a double difficulty: if cloning is (for economic reasons) restricted to the privileged few, then it will be confined to an elite demographic with far lower rates of reproduction than the rest of humanity; if, on the other hand, cloning is ultimately accessible and widespread, achieving a degree

of penetrance on the order of cellular phones, then its effect on HGV would be minimal, since there would be relative parity in its use across disparate demographics. A final possibility is that cloning could be administered in combination with GET to increase the diversity of the resulting offspring (Strong 2005).

While these questions are interesting, the focus of this paper is on GET and not cloning, largely because the potential gains from precision manipulation dwarf those associated with the crude duplication of naturally existing genomes. The notion that GET will reduce HGV turns on a critical (and highly dubitable) sociological premise: namely that individuals, when presented with the opportunity to engineer their own offspring, will tend to choose the *same or a similar set of interventions*. Some fear that this collective convergence will lead to a Brave New World of blonde haired, blue-eyed, and unhealthily proportioned people. The trouble with this idea, of course, is that it assumes there is a common conception of the good, when it is absurd to think that there is anything approaching consensus on the value and content of complex human dispositions (such as aesthetic taste, sexual attractiveness, or moral virtue). While there are certain organizing principles that are stable across cultures (e.g. morphological symmetry), they represent atolls amidst a sea of different strokes for different folks. Even if there is widespread access to

GET, disparate economic, religious, moral, political, and other cultural preferences will prevent the fixation of a small subset of phenotypes. In fact, by enabling people to act on these divergent preferences, GET could actually increase human biological diversity, allowing for new (and otherwise inaccessible) combinations of desired characteristics.[2]

Another reason to doubt that individuals and cultures will converge on a common use of GET is that the 'garden variety' is not always the best way to guarantee mating success. While there is some evidence that people are attracted to traits whose values fall close to the arithmetic mean, conformity to the morphological or behavioral status quo can also have negative reproductive consequences. A wide range of animals show an affinity for rare phenotypes in their mating decisions, a curious fact that forms the basis of an evolutionary hypothesis called 'rare male advantage,' a type of sexual selection. *Sexual selection*, which refers to differential survival and reproduction due to mate preference, can be a powerful evolutionary force, particularly in species with reduced predation pressures (such as birds and humans). Although the selection for or against a trait usually does not depend on the distribution of similar traits in the population, in *negative frequency-dependent* selection, the selective advantage of a variant is inversely proportional to its frequency. In the case of negative frequency-dependent sexual selection, this advantage is

[2] These ideas are due to a series of fruitful discussions with Allen Buchanan.

due to female mate preference for rare or minority males (Singh and Sisodia 2000). The result is a 'balancing' selection regime which maintains high levels of polymorphism in the population. Interestingly, most of the traits that are candidates for genetic enhancement are either directly or indirectly implicated in mate selection. This is not surprising, given the extraordinary ontogenetic burdens people endure in order to increase their appeal to the opposite sex, or to advance their standing among members of the same sex.

In sum, whether GET will reduce genetic diversity depends on the type of variation in question. Because the bulk of HGV is neutral, it will remain unaffected by GET, steadily accumulating variation in the absence of selection. Only the tiny fraction of functional DNA that actually matters to consumers would be subject to modification. And even if the same traits were singled out for modification, their character states would not be uniformly chosen, given that different cultures, and individuals within cultures, do not share a singular conception of the good. Finally, sexual balancing selection, global gene exchange, and human-maintained gene banks can prevent the few homogenized traits from becoming irrevocably fixed in population. For all of these reasons, it is unlikely that GET would reduce human genetic diversity to any significant extent, especially if

reproductive decisions are reserved to the individual in the private sphere, rather than mandated from the top-down by coercive political institutions.

Nevertheless, some authors contend that even small declines could have grave evolutionary consequences (Suzuki and Knudtson 1990; Lederberg 1966). This seems reasonable enough. The central issue should not be whether there is a *net* change in HGV, since an average increase in total human diversity is consistent with there being highly homogenous sub-populations which incur evolutionary costs. For the remainder of this paper, therefore, I will simply assume arguendo that GET will lead to substantial reductions in HGV, either locally, globally, or both. The question I want to address is whether this lack of biological diversity would, as some bioethicists claim, (a) increase our susceptibility to disease or (b) impair the adaptive flexibility of our species. I will show that neither scenario is plausible, let alone ineluctable.

4. Will Genetic Engineering Technology Increase Our Susceptibility to Disease?

Skeptics frequently invoke agricultural disasters in issuing bleak prognoses about the potential evolutionary consequences of genetic engineering. If the widespread cloning of potato varieties or grape vines (discussed in section 1) could result in ecological catastrophe, why should

the same lessons not apply equally to human beings? To understand why GET is unlikely to increase the susceptibility of human populations to disease, we must delve deeper into the mechanisms of biological variation and its relationship to pathogen resistance.

In sexual organisms, the two major sources of genetic variation are mutation and recombination. While the sexual combination of genomes (referred to as 'out-crossing') can generate a perpetual stream of selectable variation, doing so runs the risk of producing deleterious variants and breaking down salutary gene combinations that would otherwise go to fixation under selection. The risk was apparently worth it, however, at least for complex multicellular animals virtually all of which combine genomes instead of reproducing asexually. The ubiquity of sex presents an evolutionary paradox: why would organisms rest content with getting only half of their genes into the next generation, when asexually they could pass on *all* of them? To put it slightly differently, why should animals invest so much time, energy, and risk in mate search and copulation, only to relinquish 50% of their genome? Selection would not have countenanced such a massive cost to fitness were it not offset by some greater gain.[3]

[3] The mystery of sex surrounds not only its origins but also its maintenance. For reasons that are largely unknown, unisexual vertebrate lineages are rare and evolutionarily short-lived in the wild, despite the accessibility of parthenogenesis-conferring mutations (Adams et al. 2003).

Although the origin of sex is controversial, there are two widely received and mutually non-exclusive explanations. The first is that sex evolved to repair DNA damage from X-rays, UV light, and coding errors that could be detrimental to the phenotype (Michod and Long 1995). During the crossing-over phase of meiosis, the chromosomes align, enabling the repair of damaged portions of the genome by copying the 'correct' opposing sequences. The second explanation of sex, and the one more pertinent to the present discussion, is that recombination evolved as a means of conferring resistance to pathogens or parasites (Hamilton, Axelrod and Tanese 1990). This explanation is premised on a 'matching-alleles' model of infection genetics (Agrawal and Lively 2002), according to which an exact genetic match is required for infection (in contrast to 'universal virulence' models, wherein a pathogen can infect all host genotypes). The strategic evolutionary interaction between host and parasite leads to the so-called 'Red Queen' effect, according to which co-evolving lineages must constantly evolve in order to maintain their present fitness levels (Ridley 2003; Van Valen 1973). Anti-parasite adaptations are bound for obsolescence, particularly given the short life cycle of parasites which gives them an evolutionary rate advantage over their relatively long-lived hosts.[4]

[4] To avoid a potential cross-disciplinary confusion, note that the terms "parasite" and "parasitism" are used as *functional* concepts in evolutionary biology, where

It is widely accepted that genetic diversity (i.e. an array of genotypes) is an important factor in protecting populations from infectious agents (Spielman et al. 2004; Altizer, Harvell and Friedle 2003; Coltman et al. 1999; Meagher 1999). In the wild, in-breeding, founder effects, and habitat fragmentation can all serve to decrease gene flow between natural populations. In the context of GET, however, the fear is that pervasive genetic modification will lead to biological uniformity, rendering human populations more susceptible to pathogens. But a closer examination will reveal that it is not genetic diversity per se, but rather a *particular sort* of genetic diversity, which bears on host-parasite dynamics. The upshot is that only a minute fraction of potential genetic interventions could impact on disease resistance, and even these not incurably so.

Most studies investigating the role of variability in disease resistance have used neutral genetic markers as the metric for populational diversity. However, variability in neutral loci is only an indirect measure of the correlation between diversity and disease resistance, since it essentially serves as a proxy for variation in functionally important sequences, such as those which comprise the major histocompatibility complex ("MHC"). The MHC is a group of closely linked

---

they refer to a physically intimate and fitness-assymterical relationship between two species, and thus include organisms ranging from bacteria to the cuckoo. By contrast, in medicine and public health (including the field of "parasitology"), the term refers exclusively to *eukaryotic* parasites, and excludes viruses and bacteria.

genes in the mammalian genome responsible for immune recognition, and it is a major determinant of susceptibility to infectious and autoimmune disease. The MHC produces molecules which bind to the antigens of intra/extracellular pathogens, presenting them for appropriate T-lymphocyte response.[5] In the course of coevolution, pathogens develop novel forms of antigenicity to evade host immune recognition, and hosts, in turn, evolve new combinations of MHC genes in order to identify and destroy the immune-dodging pathogens.

Given its essential role in immune response, it should come as no surprise that the MHC cluster is the most diverse of its kind in the vertebrate clade (Robinson et al. 2003). Host organisms with more MHC alleles and allelic combinations can recognize a wider range of pathogen-derived antigens, reducing the incidence and intensity of parasitic infection (Kurtz 2003; Alberts and Ober 1993). In contrast, variability in junk DNA alone (without a corresponding diversity in functional material) is not associated with pathogen resistance (Schwensow et al. 2007; Holderegger, Kamm and Gugerli 2006).

Therefore it is not genetic variation per se, but rather *adaptive* genetic variation, which confers disease resistance on a population. To be

[5] Initially, MHC protein polymorphism may have arisen in single-celled eukaryotes in order to maintain cell membrane diversity, which can obstruct viral 'grafting,' or the passing of viral material from one host cell membrane to another (Forsdyke 1991).

even more precise, it is not adaptive variation per se, but *immuno-relevant* adaptive variation, which underwrites host resistance to pathogens. A more targeted approach to GET and cloning—one aimed at maintaining the right sort of genetic diversity—could substantially reduce the risk of infectious disease. Therefore, even if we assume that GET would narrow the range of HGV, we can significantly reduce the chances of an epidemic by deliberately preserving high levels of polymorphism in the immuno-relevant sections of the genome.

Finally, maintaining a large pool of naturally existing genetic variation may not even be a crucial asset in disease prevention and control. In contrast to other animals, and to those unfortunate individuals living prior to the germ theory of disease, contemporary human society need not sit idly by as its population is ravaged by a virulent epidemic. Unlike medieval Europeans, we are not forced to wait patiently until favorable variants have spread throughout the population, and herd immunity is achieved. To rely on HGV to see us through the coming plague would be not only epidemiologically absurd, but morally tragic. Ancestral human populations had to sustain enormous death tolls from small pox and bubonic plague in order to attain pathogen resistance. The most effective way of curtailing, containing, and ultimately eliminating an outbreak, however, is through a rapid *environmental-behavioral* response, not by

waiting for the gradual process of Darwinian evolution to run its course (a process which can take hundreds, thousands or even millions of years, depending on mutation rates, population structure, selection pressures, and the type of the adaptation in question). Canonical methods of disease control include a speedy assessment of the threat, public education on ways to prevent transmission, the provision of clean water, food and sanitized shelter, the disinfection and proper disposal of waste products, vector control, timely burials, hand-washing, quarantine, and mass vaccination (Connolly 2005).

Add to these 'low-tech' containment practices the molecular power of GET, and you have an extraordinarily capable defense against infectious disease. Unlike prophylactic measures which rely solely on environmental modulation, GET enables us to identify and synthesize the chemical functions of resistant genotypes, and to produce and distribute vaccines in the prevention and treatment of epidemics. Collectively, these methods are far more effective than natural selection in controlling an outbreak, and none are contingent on the range of HGV. Most importantly, they allow us to avoid the myriad unnecessary deaths that would inevitably occur along the winding and treacherous road to a Darwinian solution. Genetic diversity can conquer virtually any epidemic, but its victory will always be a Pyrrhic one.

While the *phylogenetic solution* is nasty, brutish and long, the eminent flexibility of human cognition and behavior offers an *ontogenetic solution* that can not only realize the same ends that natural selection is capable of achieving, but it can do so much more quickly, reliably, and with far less human carnage. GET can introduce favorable variants 'laterally' (outside of reproduction), offering a powerful mode of genetic transmission that is otherwise inaccessible to complex organisms (Powell and Buchanan, forthcoming). In this way, GET can combine and integrate variation from different human lineages, as well as genes found in other species and even those synthesized in the laboratory.

The second reason relates to human intentionality. When biologists say that variation is 'random,' they do not mean that mutation is equally likely in all directions, but rather that it is statistically unrelated to adaptation. The EHA presupposes, however, that variation is blind not only to natural selection (which it is), but also to intentional beings like ourselves (which it is not). It assumes that humans are in no better position than Mother Nature to determine which variants are fit or will be fit in the future. Despite its muddled ontology, intentionality injects a forward-looking element into the evolutionary process that the 'blind watchmaker' will never benefit from.

The argument in this section may be summed up as follows. Even if human genetic engineering reduced the range of adaptive DNA (a prospect I find unlikely for the reasons offered in section 3), there is no reason to believe that doing so would necessarily affect levels of immuno-relevant polymorphism. Because only the latter type of genetic variation affects pathogen resistance, a carefully monitored GET regime can substantially reduce the risks of human biological monoculture. At any rate, behaviorally-mediated response is a far more efficacious and morally acceptable way of dealing with an outbreak than waiting for natural selection to run its deadly course. By combining GET with established methods of disease control, we can overcome many of the physiological and moral obstacles which confront the natural origination, spread, and fixation of disease-resistant variation.

5. Will Genetic Engineering Technology Impair the Evolvability of our Species?

Even if a decrease in HGV will not render us more susceptible to disease, it is still possible that a shrinking sphere of genetic diversity could ultimately diminish the evolvability, or adaptive potential, of the human species (Suzuki and Knudtson 1989). One fear is that GET could position the human species in such precise congruity with the environment that it

becomes a hyper-specialist, unable to roll with the punches as they are thrown in the ordinary (and extraordinary) course of evolution. Another worry is that GET will operate on short-sighted motivations and flawed scientific beliefs, resulting in the elimination of potentially favorable variation. In order to evaluate these claims, we must examine the relationship between biological diversity and evolvability.

One of the central questions of macroevolution concerns the differential survival and reproduction of taxa across deep evolutionary time. Why do some groups persist for hundreds of millions of years, while others go extinct almost as quickly as they appeared? While there is no uncontroversial answer to this question, it is becoming increasingly clear that the notion of *evolvability* will be integral to any complete explanatory picture of macroevolution. Although its precise definition is contested, in the most basic sense evolvability relates to the tendency of mutations to increase the fitness of a lineage. Generally speaking, the more variation that selection has to work with, the more creative it can be in navigating the adaptive landscape (Wagner and Altenberg 1996); this in turn increases the chances that the lineage will conduct a successful evolutionary 'search' and catch the gradient of a superior fitness peak.[6] In

[6] The 'adaptive landscape,' introduced by Sewall Wright in the 1930s, is a topographic representation of the function between individual genotype/phenotypes and the environment. The fitness landscape is comprised of fitness peaks and valleys, and populations will tend to climb the nearest peak.

one sense, host-parasite co-evolution is a subset of evolvability, since it entails that the host respond to new adaptive challenges initiated by the parasite, and vice versa, in perpetuity. But above and beyond facilitating strategic maneuvers in a local evolutionary arms race, evolvability-conferring traits can, in Dawkins's words, act as "evolutionary watersheds" which open the "floodgates to future evolution" (1989, 218).

Evolvability is affected not only by the existing range of variation, but also how that variation is causally distributed. The more interdependencies there are between functional developmental systems, the more likely it is that mutations will damage the phenotype, and the less wiggle room there is for viable phenotypic variation. For this reason, evolvability depends in large part on various 'deconstraining' mechanisms that reduce the number of links between organismic processes (Raff 1996). These include (inter alia) modularity, canalization, buffering, gene duplication, and functional redundancy, all of which increase the robustness of the phenotype against microenvironmental perturbations (such as mutations or developmental noise) (Crow and Wagner 2006; Wagner and Schwenk 2000). Together, these mechanisms prevent small genetic changes from having a catastrophic effect on the phenotype.

The assumption is that if selection (and only selection) is operating on a population, mean fitness will not decrease.

Developmental robustness not only affords the phenotype with an ontogenetic margin of safety, but it also allows for the accumulation of hidden but potentially useful variation (Wagner 2003), which can subsequently be co-opted in the service of a new functional task (Kirschner and Gerhart 1998). The larger and more diverse this cache of genetic potential, the greater the adaptability of a lineage (Levenick 1999). Stephen Jay Gould referred to this stock of evolutionary potential as the 'exaptive pool' (2002, 1277). The exaptive pool is comprised of three main types of variation: (1) neutral variation which has accumulated in buffered/redundant developmental networks, (2) adaptive variation, or genes that are currently under selection but whose function can be diverted in the service of a new task, and (3) spandrels, or the non-adaptive by-products of adaptive variation. Together, these provide the necessary raw materials for future evolutionary change (Chipman 2001).

Of these three types of variation, neutral genetic evolution is arguably the most important factor in evolvability, for several reasons. First, neutral sequences make up an enormous fraction of the total gene pool. Second, genotypes that code for important functions are inextricably bound-up with the phenotype and thus effectively off-limits to directional selection. It is precisely because of their non-functionality that neutral portions of the genome are more amenable to selective cooptation. Third,

neutral evolution allows natural selection to explore a much wider range of phenotypic search space, preventing a lineage from becoming ensnared in a local optimum. By drifting around the adaptive landscape and away from its local pedestal, a lineage increases its chances of stumbling upon the gradient of a superior fitness peak (Ebner, Shackleton and Shipman 2002).

The fact that junk DNA is a vital component of the exaptive pool has important implications for the present discussion. Because consumer capital and (hence) engineering effort will not be expended in order to modify genomic sequences that have no tangible effect on the phenotype, this vast source of co-optable diversity will remain unaltered by GET. In fact, by modifying genes that mediate developmental correction mechanisms, GET could be used to significantly *increase* the levels of neutral variation and hence the evolutionary flexibility of a lineage.

But most important of all, evolvability and the co-optable HGV on which it depends may be a less important factor in the survival of our species than other sources of diversity, such as phenotypic plasticity. In contrast to evolvability, *phenotypic plasticity* is the property of an organism, not a lineage; it refers to the ability of a single genotype to generate an array of phenotypes (including behaviors). Humans are not among the most *morphologically* variable species—compare, for example, the average human family with that of the social insect colony, which

features a caste-based system of soldiers, workers, queens etc. Nor do we occupy a particularly arborescent branch of the tree of life—our lineage is maximally depauperate, as we are the only remaining species of our genus. We do, however, boast the most robust cognitive and behavioral repertoire in the history of life. We are symbol manipulators, cultural transmitters, and niche constructors par excellence. We deliberately and radically transform our selective environment, and we transmit those changes 'vertically' (to offspring) and 'horizontally' (to conspecifics). In this way, phenotypic plasticity buffers the species against environmental fluctuations, obviating or at least significantly diminishing the evolutionary 'need' for HGV.

Even more fundamentally, we must be careful not to equate either survivability or evolvability with the good, or for that matter, with each other. The fact that GET could reduce the longevity of the species is not an irrefragable or even peremptory reason for rejecting it (Powell and Buchanan, forthcoming). Everyone who travels in an automobile, plays a sport, or eats a cheeseburger recognizes that life is not simply about maximizing one's life span. Likewise, the costs associated with phylogenetic persistence may be outweighed by the gains to be had over a shorter but more agreeable span of time.

But even if we assume that the survival of the human species is an absolute moral goal, it still does not follow that evolvability is a desirable characteristic. This is because the concept of evolvability is different from, and perhaps even antipodal to, the notion of survivability. The latter refers to the tendency to persist, while the former entails the disposition to change. These two tendencies can run in tandem, but they can also come into conflict. The ability to persist may require some flexibility for future change, but there is a point at which the requisite change is so overwhelming that it may be said to negate persistence. At what moment this happens I cannot say; but there is no shame in this confession, as neither have philosophers in thousands of years been able to agree on when the famous ship of Theseus, remodeled plank by plank over Athenian generations, ceases to be the same ship. The only point I wish to make is that the disposition to evolve can in some circumstances entail the disposition to go extinct.

To understand how this could be so, one must recognize that 'extinction' in macroevolutionary terms is very different from that term as it is used in the more colloquial sense, or for purposes of moral consideration. When most people are asked to think of a 'species', they will tend to conjure the *biological* version of the concept (due to Mayr 1942), which defines the species as the most inclusive set of (potentially)

interbreeding organisms. However, many evolutionary biologists have rejected the notion that species are (or *only are*) sets of organisms with shared characteristics, in favor of a *phylogenetic species concept* which groups species according to common ancestry (Hull 1987). On this view, the same phylogenetic species at time T may be phenotypically distinct (or even wholly unrecognizable) at time T plus or minus 1, since a shared ancestry does not imply a shared set of characteristics. The upshot is this: that the human species persists in macroevolutionary terms does not imply the survival of any of the attributes that we associate with 'human nature,' or that we otherwise deem worthy of preservation. And likewise, that the human species goes extinct in the biological sense does not entail the annihilation of those characteristics we value in ourselves.

Evolvability is heavily contingent on population structure. Larger interconnected populations exhibit higher trait continuity but a lower capacity to evolve (due to gene flow which dampens founder effects). Small isolated populations with a tendency to break-off into sister or daughter species can help maintain a lineage over deep time, but it can also cause the extinction either of the parent population, or the traits traditionally associated with it. Would we consider evolvability a desirable thing if it meant a future without beings that we could even loosely call human? In an interesting twist, consider that GET could actually be used to buffer the

human species *against* its tendency to evolve, preserving the valued attributes of human nature.

If the preceding analysis is correct, then GET does not pose an unavoidable or even colorable risk to the immediate health or long-term survival of the human species. To the contrary, we should cling to genetic engineering technology much as our early ancestors cradled fire—for it may be the key to our survival in a perennially hostile world. I do not expect (nor do I desire) that the skeptical reader stop worrying and love genetic engineering technology—but I do hope that together we have the courage to think clearly about the risks and benefits of this awesome technology.

***

BIBLIOGRAPHY

Abouheif, E. "Parallelism as the pattern and process of mesoevolution." *Evolution and Development* 10, no. 1 (2008): 3–5.

Agrawal, A., and C.M. Lively. "Infection genetics: gene-for-gene versus matching alleles models and all points in between." *Evolutionary Ecology Research* 4 (2002): 79–90.

Alberts, S.C., and C. Ober. "Genetic Variability in the Major Histocompatibility Complex: A Review of Non-Pathogen-Mediated Selective Mechanisms." *Yearbook of Physical Anthropology* 36 (1993): 71-89.

Altizer S., D. Harvell, and E. Friedle. "Rapid evolutionary dynamics and disease threats to biodiversity." *Trends Ecol Evo* 18 (2003): 589-596.

Amundson, R. "Two Concepts of Constraint: Adaptationism and the Challenge from Developmental Biology." *Philosophy of Science* 61 (1994): 556-578.

Arendt, J., and D. Reznick. "Convergence and parallelism reconsidered: what have we learned about the genetics of adaptation?" *Trends in Ecology and Evolution* 23, no. 1 (2007): 26-32.

Baum, D.A., and M.J. Donoghue. "Choosing among Alternative 'Phylogenetic' Species Concepts." *Systematic Botany* 20, no. 4 (1995): 560-573.

Baylis, F., and J.S. Robert. "The Inevitability of Genetic Enhancement Technologies." *Bioethics* 18, no. 1 (2004): 1-26.

Beatty, J. "Replaying Life's Tape." *Journal of Philosophy* 103, no. 7 (2006): 336-362.

_____, "The Evolutionary Contingency Thesis." In *Concepts, Theories, and Rationality in the Biological Sciences*, edited by G. Wolters and J. Lennox, 45-81. Pennsylvania: University of Pittsburgh Press, 1995.

Bourke, A. *'The visitation of God?' The potato and the great Irish Famine*. Dublin: Lilliput Press, 1993.

Brandon, R.N. "The Principle of Drift: Biology's first law." *Journal of Philosophy* 103, No. 7 (2006): 319-335.

_____. *Adaptation and Environment*. Princeton: Princeton University Press, 1990.

Briggs, D.E.G., R.A. Fortey and M.A. Wills (1992), "Morphological Disparity in the Cambrian", *Science* 256(5064): 1670-1673.

Buchanan, A. 2008a. Genetic enhancement and moral status (unpublished manuscript)

_____ 2008b. Human enhancement and human development. *The Kennedy Institute of Ethics Journal* (forthcoming)

Budd, G.E. (2001), "Tardigrades as 'Stem-Group Arthropods': The Evidence from the Cambrian Fauna", *Zoologischer Anzeiger* 240(3-4): 265-279.

Caldwell, M.W. (2002), "From Fins to Limbs to Fins: Limb Evolution in Fossil Marine Reptiles", *American Journal of Medical Genetics* 112:236–249.

Campbell, C. 2004. *Phylloxera: How Wine Was Saved for the World*. London: HarperCollins.

Carroll, S.B.: 2005, *Endless Forms Most Beautiful: The New Science of Evo Devo*, W.W. Norton, New York.

Cassinello, J., M. Gomendio, and E.R.S. Roldan. 2001. Relationship between coefficient of inbreeding and parasite burden in endangered gazelles. *Conservation Biology* 15: 1171-1174.

Cavalli-Sforza, L.L. 1994. *The History and Geography of Human Genes.* Princeton, N.J.: Princeton Univ. Press.

Chipman, A.D. 2001. Developmental exaptation. *Evolution and Development* 3(5): 299-301.

Coltman, D.W., J.G. Pilkington, J.A. Smith, J.M. Pemberton. 1999. Parasite-mediated selection against inbred Soay sheep in a free-living, island population. *Evolution* 53: 1259-1267.

Conner, J.K. and D.L. Hartl. 2004. *A Primer to Ecological Genetics*. Massachusetts: Sinauer.

Connolly, M.A. 2005. Communicable disease control in emergencies: A field manual. World Health Organization.

Conway Morris, S. (2003), *Life's Solution: Inevitable Humans in a Lonely Universe*, Cambridge University Press, Cambridge.

_____ (1998), *The Crucible of Creation: The Burgess Shale and the Rise of Animals*, Oxford University Press, Oxford.

Crick, F.H. (1968), "The origin of the genetic code", *J. Mol. Biol.* 38(3):367-379.

Crow, K.D. and G. Wagner. 2006. What Is the Role of Genome Duplication in the Evolution of Complexity and Diversity? *Mol. Biol. Evol.* 23(5): 887–892.

Davison, E.H. and D.H. Erwin. 2006. Gene Regulatory Networks and the Evolution of Animal Body Plans. *Science* 311: 796-800;

Dawkins, R. 1989. The Evolution of Evolvability, in *Artificial Life*, C. Langton, ed. Addison-Wesley.
_____ 1982. *The Extended Phenotype* Oxford: Oxford University Press.

Dennett, D.C. (1995), *Darwin's Dangerous Idea*, Simon & Schuster, New York.

Diogo, R. (2005), "Evolutionary Convergences and Parallelisms: Their Theoretical Differences and the Difficulty of Discriminating Them in a Practical Context", *Biology and Philosophy* 20, 735-744.

Ebner, M., M. Shackleton, R. Shipman (2002). How Neutral Networks Influence Evolvability. *Complexity* 7(2):19-33.

Foley. R.: 1999, 'Pattern and Process in Hominid Evolution', in J. Bintliff, (ed.), *Structure and Contingency: Evolutionary Processes in Life and Human Society*, Leicester University Press, London.

Forsdyke, D.R. 1991. Early Evolution of MHC Polymorphism. *Journal of Theoretical Biology*. 150:451-456.

Foster, S. A. and J.A. Baker (2004), "Evolution in parallel: new insights from a classic system", *Trends in Ecology and Evolution* 19: 456-459.

Gehring, W.J. and K. Ikeo: 1999, 'Pax 6: Mastering Eye Morphogenesis and Eye Evolution', *TIG* **15(9)**, 371-377.

Gould, S.J. 2002. *The Structure of Evolutionary Theory*. Belknap.

_____(1989), *Wonderful Life: The Burgess Shale and the Nature of History*. W.W. Norton and Co., New York.

_____(1977), *Ontogeny and Phylogeny*, Harvard University Press.

Hall, B. K. (2007), "Homology or homoplasy: dichotomy or continuum?" *J. Hum. Evol.* 52: 473–479.

_____(2003), "Decent with modification: the unity underlying homology and homoplasy as seen through an analysis of development and evolution", *Biol. Rev.* 78: 409–433.

Haas, O., and G.G. Simpson: 1946, 'Analysis of Some Phylogenetic Terms with Attempts at Redefinition', *Proc. Amer. Phil. Soc.* **90**, 319-349.

Habermas, J. 2003. *The future of human nature*. Cambridge: Polity Press.
_____ 2001. An argument against human cloning: three replies. *The postnational constellation: political essays*. Cambridge: Polity Press.

Hamilton, W.D., R. Axelrod and R. Tanese. 1990. Sexual reproduction as an adaptation to resist parasites (a review). *Proc. Natl. Acad. Sci. USA*. 87: 3566-3573.

Harris, M.P., S.M. Hasso, M.W.J. Ferguson, and J.F. Fallon: 2006, 'The development of archosaurian first-generation teeth in a mutant chicken', *Current Biology* **16**, 371-377.

Holderegger, R., U. Kamm, and F. Gugerli. 2006. Adaptive vs. neutral genetic diversity: implications for landscape genetics. *Landscape Ecology* 21:797–807.

Hull, D.L. 2001. Replicators and interactors. *Science and Selection*. D.L. Hull, ed. Cambridge: Cambridge University Press.

_____1987. Genealogical actors in evolvogical roles. *Biology and Philosophy* 2: 168-184.

Hunter, J., J.P. Jernvall: 1995, The Hypocone as a Key Innovation in Mammalian Evolution', *Proc. Natl. Acad. Sci. U.S.A.* **92**, 10718-10722.

Kass, L. 2002. Cloning and the posthuman future. *Life, liberty and the defense of dignity. The challenge for bioethics*. San Francisco: Encounter.

_____ 1998. The wisdom of repugnance, L. Kass & J.Q. Wilson, eds. *The ethics of human cloning*, Washington, D.C.: AEI Press.

Kimura, M. 1983. *The Neutral Theory of Molecular Evolution*. Cambridge: Cambridge Univ. Press.

Kirschner, M. and J. Gerhart. 1998. Evolvability. *Proc. Natl. Acad. Sci. USA* 95: 8420–8427.

Kurtz, J. 2003. Sex, parasites and resistance – an evolutionary approach. *Zoology*. 106: 327–339

Land, M.F. and R.D. Fernald: 1992, 'The Evolution of Eyes', *Annu. Rev. Neurosci.* **15**, 1-29.

Lederberg, J. 1966. Experimental Genetics and Human Evolution. *American Naturalist* 100(915): 519.

Levenick, J.R. 1999. Swappers: Introns promote flexibility, diversity and invention, in *Proceedings of the Genetic and Evolutionary Computation Conference*. Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, Jakiela, V. M., Smith, R. E., eds. San Francisco: Morgan Kaufmann Publishers.

Losos, J.B., T.R. Jackman, A. Larson, K. de Queiroz, and L. Rodríguez-Schettino (1998), "Contingency and Determinism in Replicated Adaptive Radiations of Island Lizards", *Science* 279(5359): 2115-2118.

Mayr, E. 1942).*Systematics and the Origin of Species from the Viewpoint of a Zoologist*. New York: Columbia University Press.

McShea, D.W., and R.N. Brandon. "Diversity and Complexity: the zero-force expectation for evolution." Unpublished manuscript. Available from http://www.duke.edu/philosophy/bio/papers.html, accessed 26 August 2008.

Meagher, S. "Genetic diversity and Capillaria hepatica (Nematoda) prevalence in Michigan deer mouse populations." *Evolution* 53 (1999): 1318-1324.

Meyer, A. "Homology and homoplasy: the retention of genetic programmes." In *Homology*, edited by G. R. Bock and G. Cardew, 141–157. Novartis Foundation Symposium 222. Wiley, Chichester, 1999.

Michod, R. E., and A. Long. "Origin of sex for error repair. II. Rarity and extreme environments." *Theoretical Population Biology* 46 (1995): 56-81.

Millstein, R.L. "Chance and Macroevolution." *Philosophy of Science* 67 (2000): 603-624.

Nei, M. "The new mutation theory of phenotypic evolution." *Proceedings of the National Academy of Sciences* 104, no. 30 (2007): 12235–12242.

Nijhout, H.F. "On the association between genes and complex traits." *Journal of Investigatory Dermatol*ogy 8 (2003): 162-163.

Nozawa, M., Y. Kawahara, and M. Nei. "Genomic drift and copy number variation of sensory receptor genes in humans." *Proceedings of the National Academy of Sciences USA* 104 (2007): 20421–20426.

Patterson, C. "Homology in Classical and Molecular Biology." *Molecular Biological Evolution* 5 (1988): 603-625.

Philippe, N., E. Crozat, R.E. Lenski, and D. Schneider. "Evolution of global regulatory networks during a long-term experiment with Escherichia coli." *BioEssays* 29 (2007): 846–860.

Powell, R., and A. Buchanan. “Breaking evolution’s chains: The promise of enhancement by design.” In *Genetic Enhancement*, edited by J. Savulescu. Oxford: Oxford University Press, forthcoming.

President’s Council on Bioethics. *Reproduction and responsibility: the regulation of new biotechnologies* (2004). Report. Washington, D.C.

_____. *Human cloning and human dignity* (2002). Report. New York: Public Affairs.

Provine, W.B. *Sewall Wright and evolutionary biology*. Chicago: University of Chicago Press, 1986.

Raff, R.A. *The Shape of Life: Genes, Development, and the Evolution of Animal Form*. Chicago: University of Chicago Press, 1996.

Reich, D.E., S.F. Schaffner, M.J. Daly, G. McVean, J.C. Mullikin, J.M. Higgins, D.J. Richter, E.S. Lander, and D. Altshuler. “Human genome sequence variation and the influence of gene history, mutation and recombination.” *Nature Genetics* 32 (2002): 135–142.

Ridley, M. *The Red Queen: Sex and the Evolution of Human Nature*. New York: Harper Perennial, 2003.

Rifkin, J. *Algeny: A New Word--A New World*. New York: Viking, 1983.

Rosenberg, A. “How is Biological Explanation Possible?” *British Journal for the Philosophy of Science* 52, no. 4 (2001): 735-760.

Roth, V.L. “Character Replication.” In *The Character Concept in Evolutionary Biology*, edited by G.P. Wagner, 81-107. San Diego, CA: Academic Press, 2001.

Robinson, J., M.J. Waller, P. Parham, N. de Groot, R. Bontrop, L.J. Kennedy, P. Stoehr, and S.G. Marsh. “IMGT/HLA and IMGT/MHC: sequence databases for the study of the major histocompatibility complex.” *Nucleic Acids Res*earch 31 (2003): 311–314.

Salmon, W.C. *Statistical Explanation and Statistical Relevance*. Pittsburgh, PA: University of Pittsburgh Press, 1971.

Sanger, T.J., J.B. Losos, and J.J. Gibson-Brown. “A developmental staging series for the lizard genus Anolis: A new system for the integration of evolution, development, and ecology.” *Journal of Morphology* 269, no. 2 (2008): 129.

Schluter, D. “Experimental Evidence that Competition Promotes Divergence in Adaptive Radiation,” *Science* 266 (1994): 798-801.

Schwensow, N., J. Fietz, K.H. Dausmann, and S. Sommer. “Neutral versus adaptive genetic variation in parasite resistance: importance of major histocompatibility complex supertypes in a free-ranging primate.” *Heredity* 99 (2007): 265–277.

Shapiro, M. D., M.E. Marks, C.L. Peichel, B.K. Blackman, K.S. Nereng, B. Jónsson, D. Schluter, and D.M. Kingsley. “Genetic and developmental basis of evolutionary pelvic reduction in threespine sticklebacks.” *Nature* 428 (2004): 717-723.

Shapiro, M.D., M.A. Bell, and D.M. Kingsley. “Parallel genetic origins of pelvic reduction in vertebrates.” *Proceedings of the National. Academy of Science U.S.A.* 103, no. 37 (2006): 13753–13758.

Simpson, G.G. *Principles of Animal Taxonomy*. New York: Columbia University Press, 1961.

Singh, B. N., and S. Sisodia. “Frequency-dependent selection: Minority male mating advantage in Drosophila.” *Current Science* 78, no. 2 (2000): 141-150.

Sommer, S. “The importance of immune gene variability (MHC) in evolutionary ecology and conservation.” *Frontiers in Zoology* 2 (2005): 16.

Spielman, D., B.W. Brook, D.A. Briscoe, R. Frankham. “Does inbreeding and loss of genetic diversity decrease disease resistance?” *Conservation Genetics* 5 (2004): 439-448.

Strong, C. “Reproductive cloning combined with genetic modification.” *Journal of Medical Ethics* 31 (2005): 654-658.

Sterelny, K. “Another View of Life.” *Studies in the History and. Philosophy of Biological and Biomedical Science* 36 (2005): 585-593.

_____. "Explanatory Pluralism in Evolutionary Biology." *Biology and Philosophy* 11 (1996): 193-214.

Suzuki, Y., and H.F. Nijhout. "Evolution of a Polyphenism by Genetic Accommodation." *Science* 311 (2006): 650-652.

Suzuki, D.T., and P. Knudtson. *Genethics: The Ethics of Engineering Life*. Toronto: Stoddart, 1990.

_____. *Genethics: The Clash between the New Genetics and Human Values*. Cambridge: Harvard University Press, 1989.

Thattai, M., and A. van Oudenaarden. "Intrinsic noise in gene regulatory networks." *Proceedings of the National Academy of Sciences* 98 (2001): 8614.

Travisano, M., J.A. Mongold, A.F. Bennett, and R.E. Lenski. "Experimental Tests of the Roles of Adaptation, Chances, and History in Evolution." *Science* 267, no. 5194 (1995): 87-90.

Van Valen, L. "A new evolutionary law." *Evolutionary Theory* 1 (1973): 1-30.

Vermeij, G.J. "Historical contingency and the purported uniqueness of evolutionary innovations." *Proceedings of the National Academy of Sciences* 103, no. 6 (2006): 1804–1809.

_____. "Evolutionary Interaction Among Species: Selection, Escalation, and Coevolution." *Annual Review of Ecology and Systematics* 25 (1994): 219-236.

Via, S. et al. 1995. Adaptive phenotypic plasticity: consensus and controversy. *Trends in Ecology and Evolution* 10: 212-217.

Vogel, S. "Convergence as an Analytical Tool in Evaluating Design." In *The Optimization and Symmorphosis Debate*, edited by E.R. Weibel, C.R. Taylor and L. Bolis, 13-20. Cambridge: Cambridge Univ. Press, 1998.

_____. "Diversity and Convergence in the Study of Organismal Function." *Israeli Journal of Zoology* 42 (1996), 297-305.

Wagner, G. P., and K. Schwenk. “Evolutionarily Stable Configurations: functional integration and the evolution of phenotypic stability.” *Evolutionary Biology* 31 (2000): 155-217.

Wagner, G. “Evolutionary Genetics: The Nature of Hidden Genetic Variation Unveiled.” *Current Biology* 13 (2003): 958-960.

Wagner, G., and L. Altenberg. “Complex Adaptations and the Evolution of Evolvability.” *Evolution* 50, no. 3 (1996): 967-976).

Wake, D. B. “Homoplasy, homology and the problem of ‘sameness’ in biology.” In *Homology*, edited by G. R. Bock and G. Cardew, 141–157. Novartis Foundation Symposium 222. Wiley, Chichester, 1999.

_____. “Homoplasy: The Result of Natural Selection or Evidence of Design Limitation?” *Amer. Naturalist* 138 (1991): 543-567.

Weinreich, D. M., F. N.F. Delaney, M.A. DePristo, and D.L. Hartl. “Darwinian evolution can follow only very few mutational paths to fitter proteins.” *Science* 312 (2006): 111–114.

Zuker, C.S. “On the evolution of eyes: would you like it simply or compound?” *Science*. 265 (1994): 742-743.

BIOGRAPHY

Russell Powell was born in Long Island, NY on March 20, 1977. He received his B.A. in philosophy (summa cum laude, Phi Beta Kappa) from Binghamton University (1999), his Juris Doctor (cum laude) from NYU Law School (2002), and his Ph.D. in philosophy from Duke University (2008). Before returning to graduate school, Russell worked as an attorney in the New York office of Skadden, Arps LLP, where he specialized in mass tort and product liability litigation. Russell's academic interests are wide-ranging and highly interdisciplinary, ranging from the philosophy of science to political philosophy. Recent publications include:

- "Contingency and Convergence in Macroevolution: A Reply to John Beatty," *Journal of Philosophy*, Forthcoming
- "The Law and Philosophy of Preventive War: An Institution-Based Approach to Collective Self-Defense." *Australian Journal of Legal Philosophy* 32 (2007): 67-89
- "Is Convergence More than an Analogy? Homoplasy and its Implications for Macroevolutionary Predictability." *Biology and Philosophy* 22 (2007): 565-578
- "Fidelity to Constitutional Democracy and to the Rule of International Law." In *Routledge Handbook for International Law*, edited by D. Armstrong. Routledge, Forthcoming (w/ A. Buchanan)
- "Constitutional Democracy and the Rule of International Law: Are They Compatible?" *Journal of Political Philosophy*, 2008 (w/ A. Buchanan)
- "Breaking Evolution's Chains: The Promise of Enhancement by Design." In *Enhancement,* edited by Julian Savulescu. Oxford: Oxford University Press, Forthcoming (w/ A. Buchanan)
- "The Metaphysics and Ethics of Extinction." In *Ethics and Animals*, edited by T.L. Beauchamp and R.G. Frey. Oxford: Oxford University Press, Forthcoming.