

WHAT IS EVOLVABILITY?

Kim Sterelny

1 THE METAZOA AND THE VOLVOCEANS: TWO CONTRASTING FATES

For free-living unicellular protists, cell division is reproduction. A cell divides, and the two daughters go their separate ways. But very occasionally, something else happens. Somewhere between 700 and 800 million years ago, in a protist lineage closely related to the living choanoflagellates, cell division resulted in cell aggregation [King, 2004]. Instead of that division resulting in two daughter cells drifting off to their separate fates, the aggregation stayed together and shared a common fate. That lineage prospered. We can only conjecture at the reasons for its success: one possibility is simply that it was too large to engulf; its size made it safe from other protists. But though we do not know the advantage conferred by this experiment in collective life, we know its consequences. For I have just described the origins of the Metazoa: the lineage of multi-celled animals. So this aggregation was ancestor to a hugely diverse and hugely disparate clade.

The metazoans are one of the three great radiations of multi-celled life: the others are the green plants and the fungi. But these are by no means the only multi-celled lineages. There are different estimates in the literature, but no-one doubts that multicellularity has evolved repeatedly; more than a dozen times ([Bonner, 1998]; [King, 2004]). So consider the contrast between these highly disparate lineages and another and more recent experiment. The plants evolved from green algae [Graham *et al.*, 2000]. But they are not the only multi-celled descendants of green algae. About 75 million years ago another experiment in collective life began: that of the volvocaceans. The founder of this lineage lived in shallow ephemeral ponds (if the ecology of its descendants is any guide). Phosphate is the crucial limiting resource in these ponds, and there is reason to suppose that multicellular volvocaceans store phosphorous more successfully than their single-celled competitors because they can store it in the extracellular matrix that fills the spaces between their component cells.

There has been a modest radiation of multicellular descendants of these protists ([Kirk, 2001]; [Kirk, 2003]). Some multicelled lineages are halfway houses, intermediate between co-operative groups of individual cells, cells which retain independent reproductive fates, and structured collectives. There are volvocaceans like this, where cells stay together for a period before fissuring so that each component cell has its own separate shot at reproduction (as in *Chlamydomonas*).

But other volvocaceans are clearly structured collectives. In *Volvox carteri* there is a clear specialisation between somatic cells that enable individual volvocaceans to move through the water and stay at the surface (where they photosynthesise) and a much smaller number of much larger reproductive cells. True multi-celled organisms, not just co-operation between individual protists, have evolved from this protist lineage. Yet we find nothing like the complexity, the disparity, or the diversity of the metazoans. *Volvox carteri* is the upper limit of volvocacean complexity.

What explains this contrast between the metazoans and the volvocaceans? Questions of this kind can be asked at any grain. I have just contrasted the fate of two lineages that have experimented with multicellularity, but more specific versions of such questions arise within these lineages. Why, for example, is sex so developmentally plastic in many fish lineages but never in terrestrial tetrapods? Terrestrial vertebrates, unlike fish, never change sex in response to social cues. Lineages seem to differ from one another not just in their specific trajectories but also in the space of evolutionary possibility to which they have access. It just does not seem open to mammals to evolve the capacity to change sex in response to (say) skewed sex ratios, however advantageous that would be. Lineages contrast in the *actual pattern* of their evolutionary histories: some are strikingly more disparate (overall, or in particular respects) than other apparently comparable lineages. In some cases, it is plausible to suppose that these differences reflect differences in *evolutionary potential* rather than reflecting chance or selective environment.¹ There is a further temptation to think that these differences in evolutionary potential themselves have an explanation in the developmental biology of the organisms concerned: lineages differ at a time and over time in *evolvability*.

2 LIMITS ON VOLVOFACEAN DISPARITY

As noted above, multi-cellularity has evolved repeatedly, but only in a few lineages has this lead to a major radiation of multi-celled forms. My project in this paper is ask whether the idea of evolvability helps in explaining the different clade geometries of rich multi-celled lineages and relatively depauperate ones like the volvocaceans.² The volvocaceans are a good probe for investigating this problem, because we have a concrete hypothesis which explains the limits of volvox disparity. So let me begin by outlining that hypothesis. We can then ask: what do considerations of evolvability add to it?

The crucial idea is that while *Volvox carteri* has an established soma/germ-line distinction, it has established that distinction by a mechanism that limits somatic growth potential and (perhaps) the potential evolution of new cell types.

¹Thus no marine vertebrate has hands: but in contrast to sexual plasticity, this contrast with terrestrial life seems to have a selective explanation.

²The mundane answer, of course, is time: we should remember that animal disparity may well have been unimpressive after the first seventy five million years of metazoan evolution.

In *Volvox*, the founding cell grows before cell division. The series of cell divisions that form the adult organism then largely take place without further growth. The initial process of cell division produces an embryo of 32 symmetrical cells. But at the next division half of the cells divide asymmetrically, producing a population of both large and small cells. The large cells from this process become the *gonadia* cells: the germ-line. The somatic ancestor cells, the smaller cells from this first asymmetric division, divide many more times and so *V. carteri* ends up with a couple of thousand small somatic cells and a handful of much larger gonadia cells. The whole show then turns inside out so that the small, somatic, flagellated cells are on the outside, and the larger immobile germ-line cells are on the inside, set in an extracellular matrix. The gonadia, alone, then grow so that each ends up roughly 1000 times the size of somatic cells. (See [Kirk, 2001] for a good discussion of these life history issues)

Single-celled green algae, the single-celled ancestor of the volvocaceans, live a two-phase existence: they have a motile growth phase followed by an immobile reproductive phase. Kirk suggests that the genes that controlled these phenotypic changes over time have been co-opted to build a soma/germ-line division at a time. The *lag*-gene when active, suppresses the development of flagella, eyespots and chemotaxis. It originally evolved to shift the single-celled ancestor from the growth phase to the reproductive phase. For no protist can divide in a flagellated form. The cellular machinery needed for flagella is also needed to control the even division of genetic material in mitosis [King, 2004]. So within gonadia cells, this gene is on. Within somatic cells, *lag* is off and *regA* is on. *RegA* is a mutation which disables chloroplast development. Cell division in *Volvox* depends on a size threshold: once somatic cells fall below this threshold, if their chloroplasts are disabled they cannot regrow to reach it again. The germ-line/soma division in *Volvox* thus depends on genetic mechanisms that ensure that somatic cells are small, and without the potential to grow. This limits division depth and hence the organism size. It also probably limits the array of somatic cell types potentially available to this lineage.

Thus *Volvox* disparity is limited because somatic cells have a finite replication potential. The germ-soma division is organised through mechanisms which produce a somatic lineage which is restricted in total size, because the number of divisions depends on the size of founding gonadia, not on the resources harvested once division begins. Moreover, the policing mechanism which prevents somatic cells reverting to germ-line forms may restrict the potential range of cell morphology of somatic cells. The morphological architectures available in this lineage are thus sharply constrained by these restrictions on cellular constituents. Evolutionary potential is limited because the *Volvox* genotype-phenotype maps links together three features of the *Volvox* developmental program: (i) establishing a soma-germ-line distinction; (ii) intrinsic limits on the generational depth of somatic cell lineages; (iii) limits on somatic cell diversity [Nedelcu and Michod, 2004]. There is no natural way to break the linkages between these features of

volvocacean developmental programs.³ One mechanism that has an important role in policing outlaws — the establishment of the germ-soma distinction — also constrains somatic cell diversity.

3 FITNESS: A MODEL FOR EVOLVABILITY?

How does this picture of the regulation of cell differentiation relate to evolvability and to the idea that evolvability explains disparity and its limits? For surely we have just explained those limits without mentioning evolvability. There are two ways we might respond to this challenge, and the availability of these two responses exposes an ambiguity in the literature on evolvability. One idea is that the gene regulation explanation of *Volvox* disparity is shallow. The crucial issue is not whether there are *currently* mechanisms that constraint the range of cell architectures available in that lineage, but whether those mechanisms are themselves entrenched. Evolution depends on variation, and hence on the mechanisms which generate variation. Questions of evolvability are questions about the selectable phenotypic variation those mechanisms generate, and especially about how stable those mechanisms are over evolutionary change. To what extent do they themselves evolve, increasing or decreasing the variation that is available to selection?

Recent literature on bacterial evolvability has conceived of evolvability as a character that can itself evolve. In this literature, the focus has been on the evolution of mutation rates; in particular, whether elevated mutation rates show that there has been selection for increased mutation itself, or merely declining investment in error correction in impoverished environments.⁴ For while selection seems to favour elevated rates of mutation (and/or recombination) in stressful conditions that might just show that in tough times the price of more accurate replication is unaffordable. An increased mutation rate would then be an unfortunate side-effect of a stress-resistant thrifty phenotype. The case for thinking that selection might actually favour increased mutation rate is strengthened by the discovery of facultative mechanisms that increase mutation rates only in stressful conditions and then return to higher fidelity replication. These are inducible mutators: mutations which disrupt error-correction mechanisms. The case for treating these mutators as adaptations is stronger still if elevated mutation rates are targeted to specific regions of bacterial genomes. Radman, Matic and Taddei claim there are such systems. When microsatellite runs (repeated sequences of the same base pairs) are included within gene, the rate of initial copying errors increases. Normally, these errors are unimportant, for they are reliably corrected when repair mechanisms are working with their normal fidelity. But when these mechanisms are turned off, these microsatellite infested genes, and only these genes, become hypermutable

³Especially, perhaps, because *Volvox* are haploid in their multi-celled life phase and hence are vulnerable to genetic noise: mutations are not masked by working copies of the mutated allele.

⁴See [Radman *et al.*, 1999] and [Earl and Deem, 2004]; for an overview, see [Chicurel, 2001].

[Radman *et al.*, 1999]. By restricting increased variation to these genes, the cost of an elevated mutation rate is reduced.

All other things being equal, increases in the supply of variation increase the evolutionary response of a lineage to selection. Mutation rate is a component of variability. But even if selection favours elevated mutation rates, and mutation rates correlate with evolvability, this would not show evolvability to be a selectable character. My main reason for this claim is the idea that evolvability depends critically on population structure and environment. It is a characteristic of lineages, not organisms, for population structure concerns the division of a species into groups, and the flow of genes between those groups. I shall defend the idea that population structure is crucial to evolvability shortly. However, it is also true that the bacterial evolvability literature does not clearly distinguish *evolvability* from *phenotypic plasticity*. Evolvability is a property of a lineage. Plasticity is variously thought of as a property of an individual organism or a genome. A plastic genome maps onto different phenotypes in different environments. A plastic organism develops a different phenotype in different environments. For us, but not for bacteria, these two pictures of plasticity are different. For in contrast to lumbering macrobes like ourselves, bacterial genomes and mechanisms of gene expression are not fixed over the life of individual organisms: lateral gene transfer is common in bacteria and decoupled from reproduction. Gene-changing might be a mechanisms of phenotypic plasticity, not of offspring variability.

If we are inclined to accept Janzen's argument that clones should be thought of as a single evolutionary individual, the distinction between evolvability and phenotypic plasticity becomes especially problematic [Janzen, 1977]. Dawkins argued convincingly against this view in the final chapter of [Dawkins, 1982]. But that argument applies only to multi-celled clonelines, for it depends on the importance of a single-celled bottleneck through which reproduction is channelled. The existence of that bottleneck allowed Dawkins to draw an important distinction between somatic and germ-line mutations, even within clonally reproducing organisms. No such distinction can be drawn for bacteria, and hence the status of a cloneline is unresolved. But even if we think of the single cell as the individual organism, the literature on bacterial evolvability conflates it with plasticity. For example, Partridge and Barton argue that a yeast prion is a replicator that enhances yeast evolvability. It does so by changing the way messenger RNA is read. The normal form of this protein binds to the stop codon, causing translation to cease. But the variant, prion-form often fails to bind. It thus results in the translation of messenger RNA into longer protein sequences. This normally depresses the fitness of yeast with the prion. But when conditions are bad, yeast with the variant, prion-form protein sometimes do better, thus increasing the frequency of this form of protein in the yeast population. There is selection for the prion-carrying yeasts in uncertain and fluctuating environments [Partridge and Barton, 2000]. But though Partridge and Barton call this selection for evolvability, it is actually selection for phenotypic plasticity. Prion-carrying yeast alter their own phenotypes, not just the downstream phenotypes of their daughter cells. Yeast that cannot anticipate

their own environment are individually better off, on average, if they carry the prion-protein. While most of them will do marginally worse than the wild-type, a few will do much better.

In my view, evolvability is not a selectable trait, even though mutation rate may be. There is an alternative and better way of thinking of evolvability. Think of it as akin to fitness. Fitness is not a trait of an organism; the fitness of a centipede is not like its segmentation pattern or leg number. This view of fitness is not completely uncontroversial. For fitness has been understood as an explanatory property: as a measure of the congruence between an organism and its circumstances. The idea is that variations in congruence explain variations in reproductive success, both within and across populations. But it has proved difficult to specify a general congruence relation, for it would have to abstract away from the specific details of both an organism and its world. Moreover the fitness of an organism is sensitive not just to the environment but to population structure. Fitness is not just a relationship between organism and environment. The literature on the evolution of co-operation has made this clear: the fitness that is relevant to evolutionary dynamics is relative fitness, not absolute fitness. Defecting traits which lower the absolute fitness of every agent in the population can invade, so long as their effect on those without the trait is more severe than the effect on those with it. Moreover since the fitness implications of particular patterns of social behaviour often depend both on the role of the agent within the group of which it is a part, and on the role of the group within the population. Population structure helps explain relative fitness [Kerr and Godfrey-Smith, 2002]. Thus fitness is now typically understood as a dispositional property of organisms. Fitter members of a population are disposed to have more reproductive success than their rivals. But if we want to explain this variance, we appeal to the specific features of the organisms' phenotypes. It is these that explain success or failure.

However, even if fitness is not a character state that explains success or failure, it does not follow that there is nothing interesting to say in the language of fitness. There is an intermediate level of generality between the ascription of fitness differences to organisms and the analysis of specific phenotype/environment couplings. Thus within evolutionary biology we distinguish between the contribution of natural and sexual selection to organism success. In understanding the evolution of co-operative behaviour, it is crucial to distinguish between fitness effects that derive from within-group differences and those that derive from cross-group differences. We distinguish between frequency-dependent aspects of fitness and fitness deriving from optimisation. If a finch's beak is optimised to the specific seasonal conditions it encounters, the benefit it derives is insensitive to others' phenotypes. Contrast this with the female-mimicking mating strategy in the giant cuttlefish. Small males often assume the body shape and patterns of females at breeding aggregations, lurking near a breeding pair, and relying on the larger guarding male being distracted by another intruding male. When there is such a distraction, the female mimic resumes a male appearance and attempts to mate, often successfully. In contrast to the finch, this strategy depends on the supply of

intruding males. These intruders disrupt male guarding and allow female mimics their opportunity to mate [Norman *et al.*, 1999]. Sexual selection, population-structured selection, and frequency-dependent selection are conceptual tools that mark important aspects of the explanation of success and failure. Yet they abstract away from specific features of organisms and their environments.

4 EVOLVABILITY, INDIVIDUALS AND ENVIRONMENTS

I suggest that fitness understood this way is a model for understanding evolvability. Just as fitness is a dispositional feature of individual organisms, evolvability is a dispositional feature of lineages. It is not a character of an organism or its developmental system; it is not a trait. It is true that the evolvability of a lineage depends on the developmental programs of the individual organisms in the lineage. Moreover most of the literature on evolvability has focused on internal factors — features of individual developmental systems — and their role in either generating or constraining variation on which selection acts.⁵ This is certainly true of Kirschner's and Gerhart's review article that returned evolvability to the agenda of evolutionary biology [Kirschner and Gerhart, 1998]. This review was about variation, and in particular mechanisms that reduced the cost of variation by making them less likely to be catastrophic. For example, Kirschner and Gerhart bring out the crucial connection between evolvability and phenotypic plasticity. Without plasticity, evolutionary change would require impossibly co-ordinated genetic changes. Think, for example, of the evolution of new cell morphologies. Mitosis is impossible unless microtubes connect to the chromosomes in the cell and mediate their symmetrical segregation to the spindle poles at each end of the cell. Microtubes find chromosomes by mechanisms of undirected variation and selective retention: the tube is unstable and is reabsorbed unless it happens to connect with a chromosome, in which case it stabilises. Through this mechanism, the process of chromosome sorting does not have to be pre-programmed with information about the number of chromosomes, their location, or of the size and shape of the cell. This mechanism does not need to be genetically re-tuned if other mutations change the cell architecture [Kirschner and Gerhart, 1998, 8422]. These ideas are important: features of individual developmental systems are crucial to lineage-level evolvability. Indeed, in my earlier work on evolvability I attempted to specify the general characteristics of inheritance and development that promote lineage-level evolvability ([Sterelny, 2001]; [Sterelny, 2004]).

However, it is crucial to see that the evolvability of a lineage depends on much more than the developmental programs of individual organisms within the population, and the extent to which those programs vary.

1. The environment matters: for variation in the environment determines the

⁵The exception is the literature on non-genetic inheritance and its role in evolution. That literature has been much more sensitive to the role of the environment in the reliable transmission of parent-offspring similarities ([Avital and Jablonka, 2000]; [Odling-Smee *et al.*, 2003]).

fraction of the reaction norm that is expressed and thus exposed to selection. More uniform environments allow cryptic genetic variation to survive unexpressed; variation that may become important if the environment changes. This point has been developed in an important recent paper by Suzanna Rutherford. Populations have unexpressed genetic variability: for example, in natural *Drosophila melanogaster* populations, she argues, there are on average hundreds of thousands of base pair differences between the average haploid genotype. Yet these are strikingly phenotypically uniform populations. Much genetic variation is effectively neutral because it does not give rise to phenotypic variation. Yet though this variation is cryptic while the environment is stable, it can be unmasked. It is *unexpressed* difference, not *inexpressible* difference. Consider, for example, the two forms of the human *CYP1A1* gene. In non-smokers, these two forms are phenotypically equivalent. But they are associated with marked difference in lung cancer risk for non-smokers. In particular, one form of the gene makes moderate smoking much more dangerous than it would otherwise be. Likewise, new genetic variation can cause previously silent differences to be expressed. A mutant form of the heat-shock protein *Hsp90* in *Drosophila* unmasks mutations in other genes which would otherwise be silent [Rutherford, 2000].

Factors which make the environment to which a lineage is exposed more uniform are therefore important, for they allow genetic variation to be stored. In masking existing variation, they limit microevolutionary response to local variation but thereby enhance long-run evolvability by preserving genetic variation which would otherwise be eliminated from the gene pool. Moreover, there are such masking mechanisms: for organisms often have just these effects on their own environment. Often organisms in part engineer the developmental environment of the next generation, and thus make them more homogenous than they would otherwise be. Mistletoe seeds germinate only after passing through the digestive system of the mistletoe bird; the seeds of many Australian plants germinate only after fires to which their parents contribute [Odling-Smee *et al.*, 2003].

2. Population structure is relevant to the distribution of the genetic resources of the species. This is most vividly illustrated by prokaryote populations. For though prokaryotes have limited chromosomal evolution (their chromosome is circular, so there is no recombination), there is rich horizontal transfer of ready-made genetic material. Plasmids, phage DNA and transposons are all mechanisms of horizontal gene movement, with different size packets. Given the ubiquity of horizontal gene transfer, the richness of local genetic resources is obviously important [Carroll, 2002]. But though this is a very vivid case where evolutionary response depends on resources available in the local population, the general moral applies to the more familiar world of the macrobes. Microevolutionary change takes place within local populations, and if these are isolated from one another, there may well be potentially

important gene combinations which are unavailable, because the variants that would form the combination have arisen in different populations.

3. Population structure is relevant to the type of microevolutionary changes that are likely. It is a truism of evolutionary theory that microevolutionary change in a species divided into many small, isolated populations is likely to be quite different from microevolutionary change in a species divided into a few large ones. Drift and founder effects are likely to be important in the first case but not the second. Population structure is also crucial to multi-level selection. For selection at the level of groups can be important only if a species is divided into groups that vary one from another in ways important to their productivity. Multi-level selection probably plays a crucial role in the evolution of new levels of biological organization, and hence is of critical importance in expanding the space of evolutionary possibility ([Michod, 1999]; [Kerr and Godfrey-Smith, 2002]).
4. Population structure is relevant to the extent to which microevolutionary change is buffered or protected. As Eldredge has repeatedly emphasised, local adaptation in local populations is fragile; it is easily lost if demic structure is unstable, for migration from other populations breaks up gene combinations co-adapted to local circumstances. A population can stably adapt to its local world only if it extrinsically or intrinsically isolated [Eldredge, 2003].

The crucial point to remember is that significant evolutionary change typically depends on the accumulation of smaller changes. Accumulation, in turn, depends on the existence of ratchets that protect each small advance. There is an internal ratchet: high fidelity replication preserves rare favourable variations. But while internal, organism-level ratchets are necessary for evolvability, they are not sufficient. Potentially favourable change must be *preserved* by individual level inheritance processes. But these must then be *amplified* in the population and then protected (from being swamped by immigration) by population level processes, thus making a further mutational change and a further iteration far more likely. Selection is not just a passive consumer of variation: it is a creator of variation. When a phenotypic change from P to P* depends on a sequence of genetic changes from $G_0 \Rightarrow G_1 \Rightarrow G_2$, selection can make G_2 , vastly more probable by amplifying the frequency of G_1 . This makes the features of environment and population structure that are relevant to cumulative selection equally relevant to evolvability. Evolvability is not a characteristic of individuals: it emerges from an interplay between (i) individuals and their developmental systems, (ii) the populations and lineages of which they are a part, and (iii) the environments within which they are embedded.

Components of Evolvability If evolvability is a dispositional feature of a lineage, there is an important sense in which differences in evolvability do not explain differences in lineage disparity. To say that the basal metazoan lineage was highly evolvable does not explain the evolutionary radiation of the animals; likewise, to

say that the volvocaceans have very limited evolvability explains nothing about its limited disparity. But just as there is middle-level theory of the source *s* and consequences of fitness differences, there may be a middle-level theory of evolvability: characteristics shared by many lineages, and which play similar roles in explaining dispositions to evolve over time. The right way to understand proposals about evolvability is to see them as proposals about such components of evolvability: proposals about the evolution or effects of such components. I will illustrate this idea by discussing two major foci of the evolvability literature. One is the explanation of continuing evolutionary plasticity over time. William Wimsatt has pointed out that early aspects of development, in particular, should be increasingly hard to alter. For they are causally relevant to many downstream developmental stages; and the more downstream consequences a particular mechanism has, the more likely it is that changing that mechanism will have some disastrous consequence [Wimsatt and Schank, 1988]. So we should expect ontogenies to be developmentally entrenched, except perhaps for their terminal stages. And yet they are not [Raff, 1996]. The second is the mirror-image of continued flexibility: the existence and explanation of phenotypic gaps. In discussing these familiar ideas, I hope to illustrate two basic points. First: just as with fitness, we have developed conceptual tools that mark important aspects of evolutionary flexibility and inflexibility; tools that abstract away from specific features of organisms and their environments. Second: there is a serious imbalance in this work. For despite the importance to evolvability of environment and population structure, this discussion is almost entirely focused on internal aspects of evolvability.

Let me begin with a much-discussed example: modularity. A developmental module is a system that develops relatively independently of other such systems. A group of genes make up a genetic module, if there are many interactions between them and their products, and relatively few interactions between these genes and others. Developmental modules are also evolutionary modules [Brandon, 1999]: if a trait (say, tooth shape) *develops* independently of other traits, variation in that trait will be independent of variation in other traits. And if that is true, tooth shape can respond to selection without thereby altering other aspects of the phenotype. If the evolution of tooth shape is decoupled from the evolution of other traits, tooth shape will be more responsive to change in selective regimes. For it will be able to change even when there is stabilising selection on other aspects of the phenotype. The idea there is an important connection between modularity and evolvability dates to a classic paper of Lewontin, in which he argued that lineages could respond to selection only to the extent that their characters were “quasi-independent”. Tooth shape and tooth number are quasi-independent only if each can vary independently of the other. If instead, their evolutionary fate is coupled, it will be impossible to select for (say) high-crowned teeth without changes in tooth number [Lewontin, 1978]. In such circumstances, it is not sufficient for high crowns to be selectively advantageous. Such teeth will evolve only if this change is worth the price in fitness of reduced tooth number.

The point that evolutionary response to environmental change is sensitive to

linkage between traits has not been controversial, though it is difficult to test empirically. For the extent to which traits develop independently of one another is not an obvious feature of organism phenotypes. Comparison between sister groups is one way to show that traits do, or do not, vary independently of one another. Finlay and her colleagues have argued that mammalian brain structures do not show much modularity: the olfactory bulb seems to be able to shrink or grow independently of size change in other regions, but this is the exception rather than the rule [Finlay *et al.*, 2001]. In a more explicit test of the link between modularity and evolvability, Yang compared two insect sister groups: the hemimetabolous insects and the holometabolous insects. Holometabolous insects undergo full metamorphosis, and hence the morphology of the larvae is decoupled from the morphology of the adult into which the larvae eventually develops, and hence are often both morphologically and ecologically quite different from those adult forms. In contrast, the nymphs of hemimetabolous insects are quite similar to the adults they will become. It is then perhaps no coincidence that the holometabolous insects are vastly the more species rich of the two clades. This difference in diversity makes sense in the light of the developmental difference. In the more diverse clade, the morphology of larvae is decoupled from that of the adult, and this allows the larvae to differentiate ecologically and morphologically independently of the adults: they can thus avoid competing with their own adult forms, and their adaptation to their own life ways is not constrained by adult adaptation to adult lifeways.

So there has been some attempt to test the link between modularity and evolvability. But there has been more focus on the extent to which developmental programs are modular, and the ways in which modularity changes over time. Gunter Wagner and his co-workers are responsible for developing models of modularity. In these models, selection will reduce epistatic linkages between two sets of genes and their associated traits when there is directional selection on one trait and stabilising selection on the other. Such a regime will select for modifier genes which suppress epistatic connections between the stabilised trait and the evolving trait. Thus we get selection for modularity whenever genes have pleiotropic effects with opposite fitness values. This may well be quite common ([Wagner and Altenberg, 1996]; [Wagner *et al.*, 1997]). Somewhat less obviously Wagner has pointed out that modularity can increase as a side-effect of selection for genetic canalisation. The more genetic inputs to trait T, the more opportunities there are for T to be perturbed by genetic and developmental noise. So the development of a trait can be canalised by making its development sensitive to fewer genetic inputs [Wagner *et al.*, 2005].

So modularity reduces developmental entrenchment: selection regimes of the kinds Wagner characterises preserves and even increases the extent to which traits are quasi-independent, and hence help explain how it is that phenotype space has been extensively explored. Thus Roger Thomas has defined a seven-dimensional skeleton space, a space that specifies the array of possible skeleton types. In his view, this space is explored richly and quickly, because the elements of skeleton-design are module-like in their structure: for example, the materials from which

these skeletal elements can be built are either rigid or flexible, and this design decision varies independently of another: whether the distinct elements are physically unconnected, jointed or fused. The result of modular design is rapid and rich exploration. As he counts them, the Metazoa, living and extinct, have employed about 180 skeleton designs; 146 of these are found in the Burgess shale metazoans [Thomas, 2005]. If we concentrate on examples like this, we will think that the explanatory agenda is that of explaining the richness of evolutionary possibilities open to lineages.

In contrast, the developmental constraints literature has an opposite focus: that of explaining surprising limits on disparity within a lineage. In a series of recent works, Wallace Arthur has developed a particularly clear and thoughtful version of this conception of the relevance of development to evolution, with some striking examples of odd gaps in nature: gaps that seem unlikely to have a selective explanation. But though I begin with gaps, his is not just a theory of why some regions of phenotype space are empty. Consider, for example, the missing centipedes [Arthur, 2000, 54]. Geophilomorphian centipedes are a family of a thousand or so species, and there is a lot of variation in the number of segments into which their trunks are divided: these centipedes have from 29 to 191 segments in their trunks. But segment patterns come only in odd numbers: so where are the geophilomorphians with, say, trunks of 66 segments? The Lithobiomorphian centipedes are even more restricted: there are a thousand or so species of such centipedes, and they all have trunks divided into fifteen segments. In this family, there is no extant variation at all. It is hard to imagine an environmental factor that would penalise variation from these patterns, so it is natural to suspect that limits in the supply of variation explain such holes in the fabric of disparity. In the light of such examples, Arthur thinks that developmental bias — features of development which make some changes from current phenotypes probable, others possible but improbable, and other impossible, are important in explaining the direction of evolution. The supply of variation plays an essential role in determining the evolutionary trajectories, and not just be constraining the space of phenotypic possibility in unexpected ways.

No-one denies, of course, that a supply of variation is essential to evolutionary change. But it does not follow that the supply of variation makes a difference to the direction of change. If variation is typically densely and isotropically clustered around current phenotypes, then change over time will depend only on the selective environment, not the supply of variation. For variation would be available for whatever change selection favoured. We would expect variation to be isotropic if the genotype-phenotype map was one-to-one, and in which small variation in the existing genes for a trait would result in small variations in that trait. On the Fisherian conception of the relationship of genotypes to phenotypes, phenotype traits depend on many small-effect genes. On that model it might be reasonable to take the supply of variation to be irrelevant to the direction of phenotype change. But within population genetics it is no longer a standard to suppose that evolutionary change depends only on small-effect genes (see, for example,

[Leroi, 2000] and [Orr, 2000]). Moreover, we have good reason to suspect that the Fisherian model is not a good general account of the genotype-phenotype relation. Continuously variable traits might show Fisherian patterns of variation, with standing variation in the population generating normal variation on each side of the current mean. Beak length in Darwin's finches may well be influenced by a large number of small-effect genes. Genotypes close to actual genotypes would produce beaks differing slightly in depth, length and breadth from actual beaks. If that is right, there is a dense cluster of possible beaks varying slightly from actual beaks, and so selection can find a path through beak space in any direction by small increments. But many traits are not continuously variable in this way: changes in segment number from fifteen to sixteen cannot be understood as an incremental path through phenotype space. The same is true of other discontinuous changes: for example, the shift from the ancestral bilaterian symmetry pattern to the 5-fold symmetry of the echinoderms.

In short, there is no *a priori* reason to suppose that the supply of variation is typically unbiased. Arthur argues that the supply of variation is indeed importantly structured. In part, this is because he thinks there are failures of quasi-independence. Consider mammal leg length. Most mammals are close to symmetry, with front and back legs of near-equal length. Mammals with small front legs have small back legs, and so on as we move from shrew and mouse to dog to deer to giraffe. The pattern has exceptions: for example, kangaroos and their relatives have much longer and more robust hind legs than front legs. But there are no exceptions that run the other way: powerful forelegs and relatively dwarfed rear legs (though the Tasmanian devil comes close). So the pattern of rough equivalence between front and back looks quite robust. Perhaps it is just the result of selection: after all, it is not easy to think of ways of making a living that select for long front legs and short back ones. Arthur's alternative hypothesis is that this pattern is the result of difficult-to-alter features of mammalian development programs. Adaptive complexes like that of the kangaroo are difficult to reach because of structure in the supply of variation: the natural supply of variation will not provide much variance in length between the two limb pairs, and hence adaptive peaks with unequal lengths are quite often too distant from current variational pools for them to be available. Crucially, though, Arthur's argument is not just one which appeals to incomplete modularity to explain missing phenotypes. In his view, some phenotypes are larger genetic targets, and hence are more likely to appear, than others. Some phenotypes are multiply realisable — different genetic changes will give rise to the same phenotypic outcome. Other phenotypes have a single genetic profile; their development depends on a specific set of genes. Those phenotypic variants that are reachable by a number of tracks through gene space are more probable than those that depend on a specific sequence of gene changes. So bias in the supply of variation need not depend on biases in mutation itself. The many/many nature of the genotype \Rightarrow phenotype map induces its own biases in variation.

Suppose that there are indeed biases in the structure of variation: Arthur ar-

gues that those biases matter, because evolutionary change is path-dependent. Populations often find themselves on flat fitness landscapes near to but not on the slopes of incompatible local optima. Such populations will not evolve under selection until new variation extends the standing variation in the population to the foothills of one of the local optima. Once that variant has established in the population, selection will drive the population to the top of that local fitness peak, and the other will no longer be accessible to that population. The direction of evolution will depend on the order in which variations pop up in the population. In essence, the trajectory of a lineage through phenotype space will depend partly on the supply of variation if, first, the supply of variation is not isotropic around whatever the current phenotype happens to be; and, second, phenotypic evolution is path-dependent. Biases in the supply of variation make some variants available for selection sooner, and others later. Order matters: selective response to the first variant shifts the population⁶ in an adaptive landscape. While there would have been selection in favour of the second variant, had it arrived at the same time as the first, it is selectively penalised once evolution has shifted the population.

This sketch of modularity and of structured variation is not meant to be definitive. There are significant and unresolved problems with these ideas. One is the striking neglect of environmental contributions to evolvability; even when such contributions might help explain a problem with which the theorist is struggling. Consider, for example, Mark Ridley's interesting and important analysis of the evolution of complexity. Ridley regards complexity as a challenge for evolvability: for it seems that a lineage must already be complex before it can evolve complexity. Building complex phenotypes is information hungry. If that information is in the organism's genome, it follows that complex organisms require long genomes. But without complex policing and error correcting mechanisms (whose construction themselves require long genomes) long genomes tend to decay. Ridley concludes that complex life is "difficult" to evolve⁷ [Ridley, 2000]. An obvious, but unexplored option is to consider the role the environment might play in buffering similarity across the generations. A second problem is the reliance in the discussion of both modularity and structured variation on the concept of a genotype \Rightarrow phenotype map. That notion is problematic, and particularly so in this context. The metaphor of a genotype \Rightarrow phenotype map understates the importance of non-genetic developmental resources to development and the context sensitivity of the effects of genes on phenotypes [Oyama *et al.*, 2001]. Those problems are important, but the idea of such a map is even more problematic in the context of evolutionary transitions in individuality. In thinking of, for example, the volvocaceans, is the phenotype a phenotype of individual cells or of a collective? However, despite these problems, the discussion of modularity and of structured variation illustrates the key idea of this section. In discussing the components of

⁶and the spread of variation around new phenotypic means: variants which were once quite probable become much less likely

⁷Though given his argument, it seems that he ought to have concluded that it is impossible to evolve

evolvability, we can profitably idealise way from the specific developmental and environmental contexts of specific lineages. The constraints on volvocacean morphological disparity are exemplified in different forms in other lineages as well.⁸

BIBLIOGRAPHY

- [Arthur, 2000] W. Arthur. The concept of developmental reprogramming and the quest for an inclusive theory of evolutionary mechanisms. *Evolution and Development*, 2(1): 49–57, 2000.
- [Avital and Jablonka, 2000] E. Avital and E. Jablonka. *Animal Traditions: Behavioural Inheritance in Evolution*. Cambridge, Cambridge University Press, 2000.
- [Bonner, 1998] J. T. Bonner. Origins of multicellularity. *Integrative Biology*, 1(1): 27–36, 1998.
- [Brandon, 1999] R. N. Brandon. The units of selection revisited: The modules of selection. *Biology and Philosophy*, 14: 167–180, 1999.
- [Carroll, 2002] R. L. Carroll. Evolution of the capacity to evolve. *Journal of Evolutionary Biology*, 15: 911–921, 2002.
- [Chicurel, 2001] M. Chicurel. Can organisms speed their own evolution?. *Science*, 292(8 June 2001): 1824–1827, 2001.
- [Dawkins, 1982] R. Dawkins. *The Extended Phenotype*. Oxford, Oxford University Press, 1982.
- [Earl and Deem, 2004] D. J. Earl and M. W. Deem. Evolvability is a selectable trait. *Proceedings of the National Academy of Science*, 101(32): 11531–11536, 2004.
- [Eldredge, 2003] N. Eldredge. The sloshing bucket: How the physical realm controls evolution. In J. P. Crutchfield and P. Schuster (eds.), *Evolutionary Dynamics: Exploring The Interplay of Selection, Accident, Neutrality and Function*. Oxford, Oxford University Press, 2003.
- [Finlay et al., 2001] B. L. Finlay, R. B. Darlington, et al.. Developmental structure in brain evolution. *Behavioral and Brain Sciences*, 24: 263–308, 2001.
- [Graham et al., 2000] L. Graham, M. Cook, et al.. The origin of plants: Body plan changes contributing to a major evolutionary radiation. *Proceedings of the National Academy of Science*, 97(9): 4535–4540, 2000.
- [Janzen, 1977] D. H. Janzen. What are dandelions and aphids? *The American Naturalist*, 111: 586–589, 1977.
- [Kerr and Godfrey-Smith, 2002] B. Kerr and P. Godfrey-Smith. Individualist and multi-level perspectives on selection in structured populations. *Biology and Philosophy*, 17(4): 459–476, 2002.
- [King, 2004] N. King. The unicellular ancestry of animal development. *Developmental Cell*, 7: 313–325, 2004.
- [Kirk, 2001] D. Kirk. Germ-soma differentiation in *Volvox*. *Developmental Biology*, 238: 213–223, 2001.
- [Kirk, 2003] D. Kirk. Seeking the ultimate and proximate causes of *Volvox* multicellularity and cellular differentiation. *Integrative and Comparative Biology*, 43(2): 247–253, 2003.
- [Kirschner and Gerhart, 1998] M. Kirschner and J. Gerhart. Evolvability. *Proceedings of the National Academy of Science, USA*, 95: 8420–8427, 1998.
- [Leroi, 2000] A. M. Leroi. The scale independence of evolution. *Evolution & Development*, 2(2): 67–77, 2000.
- [Lewontin, 1978] R. C. Lewontin. Adaptation. *Scientific American*, 239: 156–169, 1978.
- [Michod, 1999] R. E. Michod. *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton, Princeton University Press, 1999.
- [Nedelcu and Michod, 2004] A. Nedelcu and R. E. Michod. Evolvability, modularity and individuality during the transition to multicellularity in volvocacean green algae. In G. Schlosser and G. P. Wagner (eds.), *Modularity in Development and Evolution*. Chicago, University of Chicago Press, pages 466–489, 2004.
- [Norman et al., 1999] M. Norman, J. Finn et al.. Female impersonation as an alternative reproductive strategy in giant cuttlefish. *Proceedings of the Royal Society, London, Series B*, 266: 1347–1349, 1999.

⁸Thanks to Brett Calcott, Patrick Forber and Mohan Matthen for their comments on earlier versions of this chapter.

- [Odling-Smee *et al.*, 2003] J. Odling-Smee, K. Laland, *et al.* *Niche Construction: The Neglected Process in Evolution*. Princeton, Princeton University Press, 2003.
- [Orr, 2000] H. A. Orr. Adaptation and the cost of complexity. *Evolution*, 54(1): 13–20, 2000.
- [Oyama *et al.*, 2001] S. Oyama, P. E. Griffiths, *et al.* (eds.). *Cycles of Contingency: Developmental Systems and Evolution*. Cambridge, Mass, MIT Press, 2001.
- [Partridge and Barton, 2000] L. Partridge and N. Barton. Evolving evolvability. *Nature*, 407(28 September): 457–458, 2000.
- [Radman *et al.*, 1999] M. Radman, I. Matic, *et al.*. Evolution of evolvability. *Annals of New York Academy of Sciences*, 870: 146–155, 1999.
- [Raff, 1996] R. Raff. *The Shape of Life: Genes, Development and the Evolution of Animal Form*. Chicago, Chicago University Press, 1996.
- [Ridley, 2000] M. Ridley. *Mendel's Demon: Gene Justice and the Complexity of Life*. London, Weidenfeld and Nicholson, 2000.
- [Rutherford, 2000] S. Rutherford. From genotype to phenotype: Buffering mechanisms and the storage of genetic information. *Bioessays*, 22(12): 1095–1105, 2000.
- [Sterelny, 2001] K. Sterelny. Niche construction, developmental systems and the extended replicator. In R. G. Gray, P and S. Oyama (eds.), *Cycles of Contingency*. Cambridge, MIT Press, pages 333–350, 2001.
- [Sterelny, 2004] K. Sterelny. Symbiosis, evolvability and modularity. In G. Schlosser and G. Wagner (eds.), *Modularity in Development and Evolution*. Chicago, University of Chicago Press, pages 490–516, 2004.
- [Thomas, 2005] R. D. Thomas. Hierarchical integration of modular structures in the evolution of animal skeletons. In W. Callebaut and D. Rasskin-Gutman (eds.), *Modularity: Understanding The Development and Evolution of Natural Complex Systems*. Cambridge, MIT Press, pages 239–258, 2005.
- [Wagner and Altenberg, 1996] G. P. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50: 967–976, 1996.
- [Wagner *et al.*, 1997] G. P. Wagner, G. Booth, *et al.* A population genetic theory of canalization. *Evolution*, 51(2): 329–347, 1997.
- [Wagner *et al.*, 2005] G. P. Wagner, J. Mezey, *et al.* Natural selection and the origin of modules. In W. Callebaut and D. Rasskin-Gutman (eds.), *Modularity: Understanding the Development and Evolution of Natural Complex Systems*. Cambridge, MIT Press, pages 33–50, 2005.
- [Wimsatt and Schank, 1988] W. C. Wimsatt and J. C. Schank. Two constraints on the evolution of complex adaptations and the means of their avoidance. In M. H. Nitecki (ed.), *Evolutionary Progress*. Chicago, University of Chicago Press, pages 231–275, 1988.