Chapter 12

Probing the Process-Based Approach to Biodiversity: Can Plasticity Lead to the Emergence of Novel Units of Biodiversity?

Davide Vecchi and Rob Mills

Abstract The history of biology has been characterised by a strong emphasis on the identification of entities (e.g., macromolecules, cells, organisms, species) as fundamental units of our classificatory system. The biological hierarchy can be divided into a series of compositional levels complementing the physical and chemical hierarchy. Given this state of affairs, it is not surprising that biodiversity studies have focused on a “holy trinity” of entities, namely genes, species and ecosystems. In this chapter, we endorse the view that a process-based approach should integrate an entity-based one. The rationale of our endorsement is that a focus on entities does not address whether biological processes have the capacity to create novel, salient units of biodiversity. This alternative focus might therefore have implications for conservation biology. In order to show the relevance of process-based approaches to biodiversity, in this chapter we shall focus on a particular process: phenotypic plasticity. Specifically, we shall describe a model of plasticity that might have implications for how we conceptualise biodiversity units. The hypothesis we want to test is whether plastic subpopulations that have enhanced evolutionary potential vis a vis non-plastic subpopulations make them amenable to evolutionarily significant units (i.e., ESU) status. An understanding of the mechanisms that influence organismic evolution, particularly when under environmental stress, may shed light on the natural “conservability” capacities of populations. We use an abstract computational model that couples plasticity and genetic mutation to investigate how plasticity processes (through the Baldwin effect) can improve the adaptability of a population

D. Vecchi (✉)

Centro de Filosofia das Ciências, Departamento de História e Filosofia das Ciências,
Faculdade de Ciências, Universidade de Lisboa, Lisboa, Portugal
e-mail: dvecchi@fc.ul.pt

R. Mills

BioISI – Biosystems & Integrative Sciences Institute, Faculdade de Ciências,
Universidade de Lisboa, Lisbon, Portugal
e-mail: rob.mills@fc.ul.pt

© The Author(s) 2019

E. Casetta et al. (eds.), *From Assessing to Conserving Biodiversity*,
History, Philosophy and Theory of the Life Sciences 24,
https://doi.org/10.1007/978-3-030-10991-2_12

247

when faced with novel environmental challenges. We find that there exist circumstances under which plasticity improves adaptability, where multi-locus fitness valleys exist that are uncrossable by non-plastic populations; and the differences in the capacity to adapt between plastic and non-plastic populations become drastic when the environment varies at a great enough rate. If plasticity such as learning provides not only within-lifetime environmental buffering, but also enhances a population's capacity to adapt to environmental changes, this would, on the one hand, vindicate a process-based approach to biodiversity and, on the other, it would suggest a need to take into account the processes generating plasticity when considering conservation efforts.

Keywords Biological hierarchy · Process-approach to conservation · Evolutionarily significant unit of conservation (ESU) · Plasticity · Baldwin effect · Evolutionary potential

12.1 Entity-Based and Process-Based Approaches Are Complementary

Biodiversity conservation poses a set of complex conceptual and practical challenges. The metaphor of healing and the analogy of nature as a patient serve the purpose of discriminating such challenges in three groups.¹ The state of nature as a patient must be diagnosed and the damage to biodiversity estimated in order to cure it via appropriate conservation actions. But in order to do so, we need to be able to take care of the patient in the right way. That is, we need to be able to characterise nature as a patient appropriately by identifying the relevant targets of treatment and conservation action. This contribution deals primarily with this conceptual challenge by arguing that proper treatment can be achieved not by exclusively targeting entities (i.e., units of biodiversity as conservation targets such as species or sites), but by concomitantly focusing on the processes generating such entities. In order to build our case, let us first explain why science requires endorsing a *complementary entity-and-process-based approach*.

Most people would identify biological entities such as organisms and species as paradigmatic. This intuitive knowledge (or folk biology) is often eventually refined, encompassing entities that can only be observed through microscopes such as cells and macromolecules. Biological knowledge is basically about what these entities do and how they develop or evolve. The focus on entities can be justified for at least two important reasons. First of all, nature can be thought in hierarchical terms as a series of part-whole relationships where the entities-wholes at a higher level of organisation are composed of entities-parts at a lower level. Secondly, the epistemological advantage of this compositional view is that an entity-based ontology can be

¹ The metaphor and analogy are developed in Casetta, Marques da Silva & Vecchi, Chap. 1, in this volume.

upheld despite ignorance of the processes leading to the generation (and, conversely, the decomposition) of wholes.²

The basic idea beneath entity-based classification is that nature is stratified into hierarchical levels of composition, i.e., “... *hierarchical divisions of stuff (paradigmatically, but not necessarily material stuff) organized by part-whole relations, in which wholes at one level function as parts at the next (and at all higher) levels*” (Wimsatt 2007, p. 201). Consider the physical hierarchy for example; the idea is that fundamental particles (e.g., fermions – quarks and leptons – and bosons) compose hadrons (neutrons and protons), which in turn compose atoms, which in turn compose molecules, which in turn compose all other solid, liquid and gaseous molecular aggregates that we can directly observe at the mesoscopic scale. Fundamental particles are the component parts of the hadron wholes; hadrons are the component parts of the atom wholes etc. Fundamental particles, hadrons, atoms etc. are entities belonging to different compositional levels. Biological hierarchies can be analogously divided into a series of compositional levels that complement the physical and chemical hierarchies: macromolecular wholes are composed of fundamental particles, hadrons, atoms, molecular parts; cellular wholes are composed of macromolecular parts; organismal wholes are composed of cellular parts etc. Thus, compositional hierarchies seem to identify “natural” and not purely human-dependent ontological components, even though the details of any ontology remain revisable in the light of scientific advances. The upshot of all this is that, if nature is indeed stratified into hierarchical compositional levels, then, on the one hand, all entities in the universe are ultimately composed of fundamental particles and, on the other, physical, chemical and biological hierarchies must be compositionally related. The corollary of this view is that compositionality implies some kind of physical reductionism because, on the one hand, everything is composed of basic physical stuff (e.g., quarks, leptons and bosons) and, on the other, the basic physical level of the hierarchy is primitive. This kind of compositional physicalism is unproblematic in many respects, even though this does not mean that chemistry and biology should straightforwardly be reduced to physics. One reason is that the explanation of the behaviour of chemical and biological systems might require reference to properties that are not ascribable to their physical components. Additionally, this kind of compositional physicalism rests on ontological fundamentalism, that is, the not so innocent assumption that we can make sense of the idea that a fundamental physical level exists at all (Schaffer 2003). This view implies the controversial hypothesis that quarks, leptons and bosons are not composed of anything at all, that they are “atoms” of composition in Democritus’s sense. However, even though compositional physicalism is somehow obvious (because all entities are merely composed of physical stuff), it is at the same time epistemologically vacuous. Let us explain why by making reference exclusively to biology.

²Of course, there exist also processes leading to the decomposition of wholes. However, the striking feature of the history of life is that it is a history of “complexification”, of generation of wholes. Thanks to Sandro Minelli for suggesting this clarification.

There are many potential reasons to argue that the compositional physicalism so far characterised provides an unsatisfactory account of biological entities. One argument is that biological entities are composed of non-physical components, a position that might be called vitalism. A somehow different argument states that biological entities possess properties that are lacked by their physical parts, a position that might be called emergentism. Here we shall focus on another kind of argument against compositional physicalism. Suppose we were to produce an inventory of all the relevant parts and their properties: could we infer the properties of wholes? No, because we would also need an understanding of the “rules of composition” governing the behaviour of parts in the production of wholes. An entity-based ontology cannot provide a satisfactory account of the properties of wholes without providing information about the nature of the processes governing the interactions between their parts. The basic point is that an entity-based ontology would provide limited knowledge of nature unless it is complemented with a process-based ontology. This is because of two reasons.

The first is that we cannot understand biological entities and their behaviour without knowledge about their structure and functional properties. Take a protein, a whole composed of a variety of amino acids with a number of physical (i.e., biophysical and biochemical) properties; the conformational properties of proteins are dependent on the properties of their component amino-acids but are not properties of these components; for instance, the specificity of proteins (i.e., their capacity to bind to a particular ligand) is given by their structure, where this structure is generated by the functional interaction of the polypeptide chains (themselves composed of amino acids) and the environment. The point is that knowledge of the physical properties of the macromolecular components of a protein is not enough in order to account for the structural and functional properties of proteins and thus to explain and predict their behaviour: composition does not account for structure and function. Thus, a static entity-based ontology consisting of an inventory of the compositional properties of parts does not exhaust the relevant biological properties of wholes.

The second reason is that the criteria for the identification of relevant biological entities must make reference to processes. For instance, consider organismal development: even if the embryo is very different from the adult capable of reproducing, they are the same biological individual. The individuality of an organism is not a property of its component parts (i.e., cells) but of the whole. Perhaps it would be better to say that it is not even a property but a process. Hennig (1966, p. 65) introduced the concept of semaphoront in order to make sense of this constrained organismal changeability. The semaphoront corresponds to the individual (e.g., a biological organism) in an infinitely small time span of its life history during which it remains unchanged. The same concept can probably be applied to any entity, physical, chemical or biological alike. But the semaphoront is a fiction, a conceptual device that should not be reified in order to vindicate a pure entity-based ontology. In this deep sense, it might be argued, biological entities should ultimately be thought in terms of dynamical processes: every supposedly static and unchangeable biological entity should be merely thought of as a portion of its life history (Dupré 2012).

At the same time, even a pure process-based ontology relinquishing any reference to entities would be epistemically useless: science also strives to classify types of entities and to uncover practically useful criteria for entity identification. This is clearly the case in biology: we want to be able to say that the embryo is the same individual as the adult, that this cell is and always will be eukaryotic during its life history, that this organism belongs to a particular species etc. Thus, in a very basic sense, an entity-based and a process-based ontology cannot but be complementary.

12.2 Entity-Based Approaches to Biodiversity Are Deficient

The intuitive allure and the epistemological advantages of entity-based compositional hierarchies are reflected in the literature on biodiversity. In fact, compositional hierarchies are prominent in biodiversity studies (Angermeier and Karr 1994, p. 691). For instance, taxonomic hierarchies stratify biological nature in terms of the part-whole relationship between species, genera, families, orders, classes, phyla, kingdoms, domains and superdomain. Ecological hierarchies stratify biological nature in terms of the part-whole relationship between populations, communities, ecosystems, landscapes, biomes and biosphere. Genetic hierarchies stratify biological nature in terms of the part-whole relationship between alleles, genes, chromosomes, genomes and pangenomes. All these hierarchies capture an aspect of biological nature (i.e., taxonomic inclusiveness, ecological nestedness, genetic organisation). To each level corresponds an entity type, i.e., a unit of biodiversity. A general problem with compositional hierarchies of the above kind is that the biodiversity units are not immaculately characterisable, in the sense that sometimes it is difficult to recognise a biological entity as an entity of a certain type, as a certain biodiversity unit. For instance, whether a population of organisms constitutes a certain species might be open to debate and might depend on which species concept we take into account³; biofilms or host-symbiont consortia might be either considered organisms or communities depending on which characterisation of organism we take into account etc.⁴ Biology does not provide clear-cut and universally-accepted criteria for the identification of a certain biological entity as an entity of a certain type because biological entities develop and evolve. But one of the relevant issues in conservation biology is whether entity-based compositional hierarchies of the above kind provide a satisfactory framework to characterise the units of conservation.⁵ In this section we shall suggest that they do not. In order to do so, we shall consider three issues. The first concerns the justification for the choice of

³See Reydon, Chap. 8, in this volume, on the debate concerning the nature of species.

⁴See Marques da Silva and Casetta, Chap. 9, in this volume, on this issue.

⁵This is just one of the many conceptual and practical challenges posed by conserving biodiversity. See Casetta, Marques da Silva & Vecchi, Chap. 1, in this volume.

hierarchy. The second pertains to the justification for the choice of biodiversity units. The third concerns the rationale for the exclusive focus on entities.

12.2.1 The Limits of Conservation Fundamentalism

The taxonomic, ecological and genetic hierarchies are somehow conflicting, even though they might be related at some level. One proposal is that they are cleanly related at the species-population-genome level because, as Angermeier and Karr (1994, p. 691) argued, “... *any population has a taxonomic identity (species), which is characterized by a distinct genome.*” This essentialist proposal is flawed at least in the sense that it is assumed that a species-specific genome exists, while what exists is a gene-pool (i.e., the totality of the genes of a given species existing at a given time, see Mayr 1970, p. 417). The species genome is thus a statistical artefact reconstructed with reference to this gene pool and ideally comprising all the genomic constituents of all genomes of present (but not past and future) organisms belonging to a species. If the species-population-genome relationship were clean, it would follow that by saving all present members of one species we would conserve all species-specific genomic variation, which is clearly not the case at least in the sense that some genomic variants have been surely lost in the course of evolution and others will be acquired. Of course, other ways of carving nature might exist and other compositional hierarchies, possibly linking the three hierarchies used so far, might be devised. Sarkar (2002) has for instance proposed that two compositional hierarchies should be used, one spatial (i.e., biological molecules, macromolecules, organelles, cells, organisms, populations, meta-populations, communities, ecosystems, biosphere) and one taxonomic (alleles, genes, genotypes, subspecies, species, genera, etc. until kingdoms, domain and superdomain). The advantage of this proposal is its parsimony, particularly the merging of the genetic and taxonomic hierarchies (which implies that genomic units – i.e., functional or structural genomic components – should be considered taxonomic ones).

Given a multiplicity of hierarchies, is there a possible justification for choosing one particular compositional one? For reasons that will be uncovered in this section, we strongly doubt it. However, let us suppose for the sake of argument that no hierarchy can be privileged. The following question is whether some units should be chosen as fundamental units of conservation. Clearly, it is practically impossible to conserve all diversity at all levels of a hierarchy, as it is practically impossible to focus conservation effort on the all-comprehensive top-level unit (i.e., biota and biosphere). How should we choose relevant units then? As a matter of fact, conservation practice seems to bypass this foundational question. As Sarkar (2005, p. 182) relates, the convention in conservation practice is to choose the “holy trinity” of genes, species and ecosystems as units of conservation. Note that these three types of entity belong to different hierarchies, however compositional hierarchies are characterised. Thus, conservation biologists seem to think that, for instance,

conserving genes is not sufficient to conserve ecosystems and vice versa. Are they correct? Let us analyse the holy trinity in detail by starting with genes.

As we have seen in the first section, compositional hierarchies betray a reductionist bias. Unlike in physics, in the life sciences this bias is not articulated as a problem concerning ontological fundamentalism (the idea that a compositional level is primitive): obviously no biological compositional level is primitive and ontologically fundamental given that all biological entities are made of physical stuff. But an analogous problem presents itself nonetheless: is there any reason to think that a particular biological compositional level is causally privileged? Usually this question is framed in terms of reduction: suppose that biological compositional level *x* is adopted as privileged, would it be possible to reduce all biological phenomena to interactions between entities at that level? Generally, the answer to this question has been negative, with few interesting exceptions, for instance in developmental biology (Rosenberg 1997; Wolpert 1994). Nonetheless, a tendency to consider the molecular level as the biologically privileged level is clearly present in many branches of biology. The reason is that it is thought that the behaviour of biological wholes should, in order to be properly understood, be unpacked in terms of molecular interactions. When we move to conservation practice, the related reductionist idea seems to be that genes are the fundamental unit of conservation because, by conserving all genomic variation, we concomitantly preserve much of the phenotypic variation that characterises the populations constituting the species and higher taxonomic levels. Sarkar (2002, p. 152) notes that this position can be justified only if some form of “global genetic reductionism” (i.e., the thesis according to which “all biological features are, in some significant way, reducible to the genes”) is vindicated. In a very clear sense, global genetic reductionism is wrong, fundamentally because phenogenesis at all levels (from transcription, translation and protein folding up to cellular differentiation and morphogenesis) is causally influenced by a variety of environmental inputs. Thus, saving all genes would not save all possible phenotypic outcomes unless we also conserved all possible developmental environments, which verges on the impossible.⁶ Interestingly, note that developmental environments (e.g., the folding environments of proteins considered in Sect. 12.1) are fundamentally ecosystems. Also note that genes are units of the taxonomic hierarchy while developmental ecosystems are units of the spatial or ecological one. This explains why a compositional hierarchy cannot be privileged over the others and why, as a matter of fact, focusing conservation efforts on units of two different hierarchies might turn out to be a necessary rather than an incoherent conservation strategy (for a similar argument, see Sarkar 2002, p. 152).

⁶There remains a possible sense in which the conservation of genomic variation goes a long way to achieve conservation of all biodiversity: if it were established that speciation (as the epitome of a lineage diversification process) completely depends on genomic change, then we would have a good argument. The issue concerns the origin of biodiversity: if it turns out that genomic change is central, then some diluted form of global genetic reductionism might be rescued in the face of phenotypic plasticity (perhaps the variation produced through plasticity would be ineffectual *per se* for speciation; see West-Eberhard 2003 for an opposite argument).

As already argued above, the idea that a peculiar genome characterises a biological species remains common in biology (despite being generally rejected in philosophy of biology). This seemingly clear link between genetic properties and species partly explains why the latter are considered an element of the trinity. After all, species are the repositories of all the genomic and phenotypic variation among its constituent organisms, where this variation is the raw material on which speciation processes work. According to Mayr (1969), species are the most fundamental unit of biological organisation. Interestingly, Mayr argued that only sexually reproducing organisms form species and that the category is not applicable to many unicellular groups of organisms (e.g., bacteria), thus betraying a bias that still characterises conservation practice too. Mayr's (1969, p. 316) argument was that species serve a specific biological function because dividing the total genetic variability of nature into discrete packages prevents the production of "disharmonious incompatible gene combinations". Conservation efforts that target species could therefore be justified as we would save all the possible "harmonious genetic combinations".⁷ However, even if we endorse the view that species are important units of biological organisation, this would not be enough to justify an exclusive focus on this biotic unit in conservation practice. One reason is that estimating biodiversity through species count is problematic.⁸ For instance, species diversity would not account for diversity at other levels of the same hierarchy. The fact that there are more terrestrial than marine species does not translate into more diversity at the next hierarchical level; in fact, there are more marine than terrestrial phyla, i.e., diversity and disparity clash (Grosberg et al. 2012); hence, by conserving an equal number of marine and terrestrial species, we might not conserve equal marine and terrestrial biodiversity at the phylum level. Conversely, species diversity would not account for genetic diversity, that is, for diversity at another level of a different hierarchy (or even of the same hierarchy if the general taxonomic hierarchy proposed by Sarkar mixing genetic and taxonomic units is endorsed); hence, for instance, by choosing to conserve indiscriminately either species S1 and S2 of genus G because one of the two is functionally redundant (in the sense that they play an equivalent ecological role in the ecosystem), we might not be able to conserve equal biodiversity at the genetic level; the reason is that one of the species might harbour more genetic diversity (its gene pool might be larger); so, supposing chimps and bonobos play equivalent ecological roles in the ecosystems, conserving bonobos with presumably much smaller gene pools than chimps (Prado-Martinez et al. 2013) would amount to failing to conserve genomic diversity.

Similar arguments apply to exclusive focus on ecosystems as the unit of conservation. This means that it is clearly difficult to justify biodiversity fundamentalism.

⁷By adding the hypothesis that genes are the most important causes of phenogenesis, we end up with the strong hypothesis that by conserving the species' characteristic gene pool (i.e., an aggregate of genomes) we are also conserving the entirety of their possible phenotypic manifestations (that is, all protein and cell types as well as all supra-cellular organismal traits), i.e., all genetic and phenotypic biodiversity.

⁸See Borda-de-Água, Chap. 5 and Crupi, Chap. 6, in this volume on this issue.

A similar position has been argued for by Angermeier and Karr (1994, p. 691). They generalise the failure of biodiversity fundamentalism by also arguing that even a focus on a single hierarchy is bound to fail, as it would lead to ignore most biodiversity. Noss (1990, p. 357) has made this point quite succinctly by arguing that “*No single level of organization (e.g., gene, population, community) is fundamental*”. Of course, the idea of taking into account 3 units of different hierarchies instead of one unit is exactly tailored to avoid such problems. But, as Sarkar (2002, p 138) has argued, “...even this catholic proposal falls afoul of the diversity of biological phenomena ...”. Thus, Sarkar argues, even avoiding biodiversity fundamentalism in some of its two forms (either focusing exclusively on a hierarchy or on a unit) would not allow accounting for “endangered biological phenomena” that are in principle amenable to conservation, such as the synchronous flowering of particular bamboo species at a distance. Sarkar argues that in order to save this peculiar phenotypic outcome, conservation efforts should neither be directed to conserve the genome of the clumps of these bamboo species, nor even conserving the species; rather, what should be conserved are the environments in which this behaviour is expressed; only by also preserving the habitats and sites where these biological phenomena occur we would be able to conserve them. Note that this argument is analogous to the one proposed above concerning the conservation of developmental environments. Developmental environments are, like habitats and sites, entities belonging to the spatial compositional hierarchy, that is, a different hierarchy than that to which genes and populations belong. We conclude that for all these reasons there is no justification for focusing exclusively on one compositional hierarchy in conservation practice. As we have showed, at least one spatial and one taxonomic unit are concomitantly needed as conservation units in order to encompass all phenotypic biodiversity (e.g., protein conformations and developmental outcomes, genetic and phenotypic variants) and all biological phenomena (e.g., synchronous flowering of bamboo). The corollary of this conclusion is that no biodiversity unit can be the fundamental unit of conservation. Rather, a variety of units are needed to encompass all biodiversity. We shall now suggest that the limits of biodiversity fundamentalism (both in its hierarchy and unit variants) and of multi-unit approaches to conservation is arguably a symptom of a more general malaise concerning entity-based approaches to conservation practice. The fundamental question is thus whether a different kind of approach should be favoured. In particular, we ask whether there exists a rationale for the exclusive focus on entities.

12.2.2 Towards an Entity and Process-Based Approach to Conservation

As we argued in the first section, one limit of compositional hierarchies pertains to their lack of structural and functional information. The problem is thus whether conservation strategies can be devised in the absence of detailed knowledge concerning the structural properties and functional interactions between the entities

constituting the compositional levels of the hierarchy. Structural hierarchies aim to represent the organisation (e.g., the topology or network of interactions) between the parts of the relevant entity-whole, while functional hierarchies map the processes governing the causal interactions between the various parts of the relevant entity-wholes. In this sense, a structural characterisation of, for instance, a cell is the topology of the network of interactions between its components parts. The structural characterisation is not merely a list of cellular components (it is not purely compositional), but it is an organised list whereby their interactions are identified. A structural characterisation is more informative than a compositional one, but is less informative than a functional one. From a functional point of view, a cell is literally an ecosystem whereby energy and matter acquired from the environment is processed internally in such a fashion as to manufacture its component parts (Luisi 2003). This means that a functional representation of a cell specifies the causal nature of the interactions between its sub-cellular components. Consider secondly that, given that functional hierarchies aim to represent the causal interactions between the elements of a hierarchy, they do not provide merely entity-based ontologies. For instance, a functional characterisation of the cell makes reference to the metabolic interactions between nutrients, constituent proteins and other macromolecules, organelles, membrane receptors etc. In this sense, it does not purely provide an entity-based ontology but also a process-based one. The upshot is that the genetic, taxonomic and ecological hierarchies for characterising biodiversity in terms of genetic organisation, taxonomic inclusiveness and ecological nestedness are, given their compositional ethos, insufficient to capture the structural and functional aspects of biodiversity (Franklin 1988). It is for this reason that compositional hierarchies should be complemented with structural and functional hierarchies, as suggested by Noss (1990, p. 359). As soon as we look at functional hierarchies, we grasp that the focus is also on processes, not merely on entities.

The crucial question is whether knowledge of functional interactions and process is necessary in order to provide a satisfactory characterisation of biodiversity and especially of the units of conservation. Consider functional interactions first. Many species are involved in complex biological relationships such as predation and pollination. Compositional hierarchies provide information concerning the relata (i.e., the entities involved in a relation) of such interactions, but this information is oblivious to process. Pollination is an ecological function that can be realised in multifarious ways by a variety of species of insects, birds, bats, snails etc. on the one hand and flowering plants on the other. Perhaps some species play a fundamental ecological role in the pollination process (as keystone species, Sarkar 2005, p. 15) and our conservation efforts should be focused on these.⁹ It is therefore clear that knowledge

⁹The concept of keystone species can be characterised in terms of ecological centrality (when a species has many functional relationships with many different species). This characterisation, however, seems to imply a lack of specialisation on the part of the species. For instance, pollination seems to be realised in large part by highly specialised species (both plants with very few pollinators and animals pollinating very few plants) which do not have, as a consequence, many functional relationships with many different species. We would argue that the ecological centrality of a keystone species depends on its specialised functional role: a keystone species would thus be one

of this ecological role might inform conservation efforts. However, knowledge of this kind is clearly not provided by compositional hierarchies. Consider processes now. Generally speaking, two types of processes governing the behaviour of the entities identified by compositional hierarchies can be identified. First of all, those leading to the differentiation of parts. Secondly, those that, given the differentiated parts, govern their combination (i.e., combinogenesis). All natural sciences are somehow concerned with understanding the nature of the processes of part differentiation and those governing their combination. Biology certainly strives to understand differentiation and combinogenetic processes: biology is both about differentiation of part-entities (e.g., production of genetic variants, new species etc.) and about the combinogenesis of whole-entities (i.e., the emergence of new biological individuals). For instance, what are the processes that govern allelic, population and species differentiation? Theories of genomic change and speciation are part and parcel of biology of course. And what are the processes that govern genome and ecosystem formation? Equally, theories concerning genome evolution and ecology are part and parcel of biology. An entity-based approach to biodiversity is thus parasitic on biological theories concerning, among others, genomic and phenotypic change as well as biological and ecological theories concerning, among others, genome evolution, phenotypic evolution, speciation and ecosystem stability, where all these theories make a reference to processes (e.g., mutation, phenotypic plasticity, predation) impinging on a variety of biological entities belonging to various levels of various compositional hierarchies. Thus, given that reference to such processes remains invisible in compositional hierarchies, they seem by their own nature epistemologically deficient. This point is particularly relevant because it influences the characterisation of the units of conservation. Does a focus on units of biodiversity, which are biological entities, make sense without a complementary focus on their maintenance and generative processes?

12.3 Does a Process-Based Approach to Biodiversity Make Sense?

We have argued so far that an entity-based approach ignores the functional relations between the elements of the hierarchy. In a nutshell, it ignores the influence of processes of differentiation and combination of parts. In conservation science, a process-based approach would shift the focus on the processes that *originate and maintain* biodiversity. As we shall relate, the shift from entities to process has been advocated by many conservation practitioners. The argument that we shall propose does not advocate a switch to exclusive focus on process. More reasonably, we suggest that a process-based approach should integrate an entity-based one (Faith

performing (almost) exclusively a particular function (e.g., pollination) for other species. See Sect. 12.3 for a clear example of keystone species (Morris et al. 2012). Thanks to Alessandro Minelli for drawing attention to this putative tension.

2016). After all, what could it mean to conserve a process? Not much. As we already argued in Sect. 12.1, if an exclusive entity-based approach to biology does not make sense, even an exclusive process-based approach does not. The reason is obvious: processes are important because they create and maintain new entities, new units of potential conservation. Thus, to use the pollination example again, the issue is not whether we should either choose the relata (e.g., the populations) or the relationship as units of conservation. We cannot think of any other sensible way of conserving relationships and processes than by conserving their relata and their actors (i.e., the entities involved in the process). As we shall explain below, the shift to process is most prominently a shift in the ways in which we characterise the units of conservation. Particularly important in the present context is Ryder's (1986) proposal to characterise conservation units as evolutionary significant units (i.e., ESUs), that is, as populations of organisms that, for historical and evolutionary reasons, play peculiar causal roles in the processes targets of conservation. From a process-based perspective, the ultimate focus of conservation practice is on entities such as ESUs (Moritz 1999, p. 223).¹⁰ Relatedly, an important issue about the characterisation of a process-based approach concerns the kind of processes that should be taken into account. Noss (1990) considers as potential targets conservation processes that are partially abiotic such as energy cycles. Noss's is an interesting suggestion. However, it should be highlighted that, again, the focus is, ultimately, inevitably on the entities that play specific causal roles in processes. For instance, in marine environments some bacteria seem to play the role of keystone species as they might exclusively perform some specific function. For example, a limited number of bacteria (e.g., of the genus *Alteromonas*) seem to process hydrogen peroxide in the ocean, performing a crucial metabolic function that benefits the incredibly large communities of the cyanobacterium *Prochlorococcus* (Morris et al. 2012). Without these bacteria, the ecosystem would probably suffer. Conservation efforts could thus be directed to save this important geochemical process, but inevitably such efforts would focus on preserving the important ecological function that *Alteromonas* bacteria play. We thus suggest that the focus should be on the processes that govern what we called entity differentiation and combinogenesis in Sect. 12.2.2, that is, most prominently the ecological and evolutionary processes that cause the origin of ESUs.

Many inter-linked themes prominent in the conservation literature explain the shift towards a complementary entity-process-based approach to conservation. This conceptual shift finds its theoretical support in the deeper integration with the evolutionary sciences and with ecology. Most generally, Norton (2001) argues that conservation science has experienced a transition from a static to a dynamic view focused on evolving systems and ecosystem processes. This interpretive hypothesis is probably supported by a shift in the characterisation of the units of conservation from static entities – e.g., species characterised essentially in terms of species-specific genetic and phenotypic features – to historical ones with peculiar historical

¹⁰ Sarkar (2002, note 15, p. 152) has argued that focus on process is aimed to conservation of biological "integrity" rather than biodiversity. However, if the focus of a process-based approach is on entities such as ESUs, it is clearly committed to biodiversity conservation.

and evolutionary capacities. This transition has been nurtured by the dissatisfaction with prominent species approaches to conservation aimed at the maximisation of number of species saved per spatial area which are, by definition, fundamentalist and entity-biased. Ultimately, the idea is that the conservation focus should be put on the evolutionary and ecological causes of biodiversity and the preservation of process rather than on their causal effects and on the preservation of pattern. Particularly important are the attempts to identify centres of evolutionary diversity with the aim of maximising evolutionary heritage on one hand (a consequence of the integration of phylogenetic analyses with conservation biology) and the focus on the evolutionary (e.g., genetic, cf. Frankel 1974) potential of populations and historical lineages. Smith et al. (1993) argue that knowledge of the ecological and evolutionary mechanisms generating genetic diversity and of the isolating mechanisms of speciation must be part and parcel of conservation practice. Conservation practices that focus on protecting species-rich sites are doomed to fail for reasons that parallel those for which the counting-species approach did. First, such focus does not necessarily provide information on the frequency of rare species, which might not occur in areas of highest species diversity (Smith et al. 1993, p. 164). Secondly, it does not necessarily provide any information on the functional role of species and on the nature of the community dynamics of the relevant ecosystems (Smith et al. 1993, p. 165). Thirdly, it neither necessarily identifies regions with peculiar evolutionary history nor identifies lineages that are phylogenetically unique (ibid.). For all these reasons, Smith et al. propose an approach to conservation that integrates ecological and molecular information. Related to the third point above, Mace et al. (2003) have suggested that, rather than directing conservation efforts to save species, these should be directed to saving independent branches of the tree of life, that is, distinctive lineages with a long and unique evolutionary history. The rationale for this conservation strategy is that phylogenetic information permits to distinguish “cradles” of diversity from “museums”. A process-based approach informed by phylogenetic information (and hence by knowledge about evolutionary history) identifies as priority conservation taxa those that display a unique evolutionary history instead of focusing efforts on conserving patterns of species richness (Mace et al. 2003, p. 1709). Along the same lines of integration of molecular data, Moritz (1999) has proposed to address conservation problems by focusing on the maintenance and restoration of those ecological and evolutionary processes that can recreate adaptive phenotypes. In order to conserve such processes, we should aim to conserve their “effectors”, i.e., the ESUs or populations with evolutionary potential in which they play causal roles. Moritz argues that molecular studies are particularly important to infer evolutionary history. Molecular information will give us details about the evolutionary relationships between the populations of conservation focus to the extent that, for instance, “... translocations among populations that historically exchanged genes would be considered, whereas human-mediated mixing of historically isolated gene pools would be discouraged.” (Moritz 1999, p. 223). This approach aims to conserve ESUs through the restoration of connectivity between isolated populations in anthropogenically fragmented ecosystems and the destruction of “genetic ghettos” (Moritz 1999, p. 224). In synthesis, a process-based

approach to conservation might be seen as proposing an integration of varieties of ecological and evolutionary information with the aim of identifying relevant ESUs. One general characterisation of ESUs that can be extrapolated from the conservation literature reviewed so far refers to populations of organisms possessing a property of conservation interest, such as a peculiar history (i.e., being a distinctive lineage) and a crucial functional role in ecosystem welfare (i.e., being a keystone species). Even though preservation seems to be, by definition, the aim of conservation biology, it is interesting to observe that Smith et al. (1993, p. 164) have argued that the aim of conservation science is "...to promote and preserve natural dynamics." What could promotion amount to? A promotion (rather than preservation) characterisation of ESUs might refer to properties of populations such as the ability to cope with environmental stress (i.e., adaptability) or an enhanced capacity to diversify into lineages with distinctive genetic and phenotypic features. In the latter two cases, it might be said that the population ESU displays "evolutionary potential" (Casetta and Marques da Silva 2015), a property that might depend either on possessing particular genomic properties or on its tendency to respond to environmental change purely phenotypically, where such properties might be important for populations' adaptability and diversification.¹¹ In the following section we shall focus on populations that display evolutionary potential in the latter sense. The hypothesis we would like to test is whether plastic populations of a species might be considered ESUs amenable to conservation. In particular, we would like to show that plastic subpopulations that have enhanced evolutionary potential vis a vis non-plastic subpopulations make them amenable to ESU status.

12.4 Can Phenotypic Plasticity Confer Evolutionary Potential?

In this section we shall thus focus on a particularly evolutionary process, i.e., phenotypic plasticity (Fitzpatrick 2012; Forsman 2015; Miner et al. 2005; Valladares et al. 2014; West-Eberhard 2003). By plasticity we refer to the ability of the organism to react to environmental inputs with an appropriate phenotypic change during embryogenesis (developmental plasticity) and further developmental stages (phenotypic plasticity). Two main types of plasticity exist: reaction norms and polyphenisms. In reaction norms the genome allows a continuous range of potential phenotypes. On the other hand, polyphenisms are discontinuous (either/or) phenotypes elicited by the environment. The essence of plasticity is that the genome does not wholly dictate the nature of the phenotypic outcome. It is reasonably straightforward to intuit about selective advantages to phenotypic plasticity: where there exist different or varying environmental conditions that are experienced, either a) by different individuals across a population, or b) by the same/each individual through its

¹¹ See Minelli, Chap. 11, this volume, for the relationship between evolutionary potential and evolvability.

lifetime, a unique phenotype (narrow reaction norm) would be less fit than plastic responses. This is, of course, provided that an appropriate phenotype can be expressed, either sensitive to environmental conditions or genetic (West-Eberhard 1986). An example of a genetic switch is the X-Y sex determination system in mammals. Some species of buttercups (e.g., *Ranunculus flammula*) exemplify polyphenism through environmental sensitivity: they develop one of two distinct leaf types, depending on whether underwater or on land (Cook and Johnson 1968). A particularly advanced form of environmental sensitivity – potentially producing continuous phenotypic responses – is learning: the capacity to *change* behaviour in particular situations, according to past life experiences (Staddon 1983).

While these phenotypic flexibilities are interesting in their own right, and indeed potential benefits of plasticity are easy to identify (notwithstanding discussion regarding what those benefits trade off against), could there be a deeper evolutionary issue here? Could phenotypic plasticity not only have proximate effects, but also impact the course of evolution? The understanding that traits produced through plasticity are not heritable goes as far back as the nineteenth century with August Weismann's experiments showing a soma/germ-line separation. And the hypothesis now commonly known as Lamarckian evolution, that traits acquired during lifetime would be passed on to further generations – e.g., the strong biceps of a blacksmith – is not considered compatible with genetic inheritance (discounting epigenetic inheritance). But there is an intriguing suggestion that phenotypic changes could influence selection in an evolving population, and thus indirectly lead to genetic encoding of formerly acquired traits (Baldwin 1896; Osborne 1896; West-Eberhard 2003). The basic notion is that the relatively rapid exploration of phenotype space via plastic response can introduce a selective gradient towards genetic specification of that phenotype, and thus the slower genetic variation can be “guided” by lifetime exploration (Hinton and Nowlan 1987). The selective landscape experienced by a plastic population is modulated by that plasticity, in comparison to the landscape experienced by non-plastic populations. But the modulation to fitness of specific genotypes does not require that the phenotypic traits discovered are heritable, i.e., it occurs without so-called “Lamarckian” inheritance. This process has become known as the Baldwin effect (a term coined by Simpson 1953). The effect depends on the existence of phenotypic plasticity (Bradshaw 1965) having already evolved but this in itself can only facilitate the first of two phases: selection among the various phenotypes expressed for those most appropriate to the present environmental conditions. The second phase, genetic assimilation, is not a necessary consequence of the existence of plasticity, nor does it depend on a reduction in the level of plasticity.¹² The Baldwin effect has been the subject of a plethora of computational studies (see e.g., Turney et al. 1996; Paenke et al. 2009; Szajder et al. 2012), following the seminal work of Hinton and Nowlan (1987). Almost all of these works considered evolution in single-peaked fitness landscapes; but in Mills and Watson (2006) we showed that, via a Baldwinian process, a learning population is able to cross a fit-

¹² Mills and Watson 2005 further discuss how canalisation, although often implicated in studies on the Baldwin effect, is not actually a necessary mechanism for the effect.

ness valley. Here we use the same model to illustrate various scenarios, including that learning is able to repeatedly guide genetic evolution in a variable environment.

12.4.1 A Model of Plasticity

We model a population of individuals each with a string of n binary variables to represent their genotype, which specifies the phenotypes that the individual will express throughout its lifetime, through a trivial (but non-deterministic) genotype–phenotype (G-P) mapping. Specifically, for each lifetime trial i , the phenotype p_i is based on the genotype with mutation-like variation applied at a rate of μ_L , independently applied at each locus. The phenotypes from the T trials are independent from each other, and can be thought of as a cloud of points surrounding the genotypically-specified location. The individuals are bestowed with a simple capacity to learn, which is facilitated through the way that fitness is calculated: during each lifetime trial, a learning individual recalls the best solution found so far, whether it is the newest phenotypic strategy, or whether it was found long ago (see Hinton and Nowlan 1987). At the end of each generation, the individuals are selected in proportion to their fitness,¹³ and reproduce asexually. During reproduction, point mutation is applied to each gene, i.e., each gene is transmitted to the offspring with a probability of $1 - \mu_G$, otherwise with probability μ_G a new random allele is drawn (note that this model does not rule out the possibility of multiple mutations but that they are uncorrelated when they occur). The population size m is constant through time. In this model there is no way for an individual to perform less lifetime exploration, i.e., there is no mechanism for canalisation (Waddington 1953). This simplification is not meant to imply that there would never be a selective advantage to such a reduction, but rather to keep the spotlight on the consequences of plasticity.

Simulation Experiment 1 We consider the evolution of a population on a simple and abstract fitness landscape, where there are two rare phenotypes p_1 and p_2 that receive high fitness and all other phenotypes are equally bad. Here, our main question is to investigate whether the form of phenotypic learning in this model is sufficient for the population to evolve across the fitness valley between the two peaks. Accordingly, the first peak/phenotype confers high fitness ($f(p_1) = H$) and the second peak confers lower fitness ($f(p_2) = L$). The environment remains like this for s generations, after which the quality of peaks switches, such that $f(p_1) = L$ and $f(p_2) = H$.

Parameters used in this experiment: $H = 100$, $L = 10$, $f(p|p \neq p_1, p \neq p_2) = 1$, $n = 16$ genes. The separation d of the two peaks is 5 bits, and the switching interval s is 50

¹³ Since all individuals experience the same number of learning trials before selection, it could be seen as selection only occurring on adult organisms; however, the model confers benefits to successful learning earlier in the lifetime, even though we do not explicitly include phenomena such as probabilistic death without reproduction.

Fig. 12.1 Fitness measures of populations evolving in a switching environment. (a) one example run; (b) averages over 50 replicates. The switches in environment are marked by the dotted vertical lines

generations. We set the mutation rate μ_G at $1/20$ and the lifetime variation rate μ_L at $2/n$, the number of trials per individual $T = 256$, and the population size m at 200.

To see what is happening in the population, we can observe the fitness over time (Fig. 12.1). Initially all organisms of the population possess the genotype specifying phenotype p_2 , and within only a few generations some individuals in the population find a high-fitness phenotype, p_1 . Any mutation that brings the genotype closer to directly specifying p_1 will be favoured since discovering the phenotype earlier in the lifetime results in higher fitness. Accordingly, such high-fitness genotypes propagate through the population, as is reflected in the rise in mean fitness. After each switch in the environment (dashed vertical lines), we see a sharp drop in fitness, reflecting the fact that the population was adapted to a previous challenge. However, phenotypic plasticity enables individuals to rapidly re-discover p_2 , which is now the highest-fitness phenotype in the environment.

Simulation Experiment 2 Rather than fixing the rate of environmental switching, here we leave this parameter s open; and to ascertain the capacity of a plastic population to cope with such environmental change we run simulation experiments for various different values of T .

From the results in Fig. 12.2 we see two different trends: (1) populations experiencing a large number of trials T can achieve high fitness, provided the environment does not change too rapidly. When the interval is very short there is insufficient time for the population to find the high peak and assimilate it genetically. Note however that the high peak is found phenotypically by some fraction of the population: with-

Fig. 12.2 Fitness of the population in the final generation before a switch. Data show the mean and standard deviation of mean fitness, across 40 replicates

out doing so, the mean fitness could not be greater than $f(L) = 10$. (2) The smaller the value of T , the less capacity the population has to adapt to the new challenge presented by the switched environment. At its extreme, with $T = 2$, the population wholly fails to adapt under any environmental switching rate tested.

The dynamics of a non-plastic population are qualitatively different. As the mutation model permits multiple loci to change simultaneously during reproduction, a multi-locus valley could, in principle, be crossed. However, the expected time this takes grows exponentially with the number of genes that must change at once. For the 5-bit valley and $\mu_G = 1/20$ (as used above), it takes a mean of over 75,000 generations, and even though a higher rate would reduce this, even an optimal rate of $\mu_G = 5/16$ takes a mean of 340 generations (mutation rates that are so high introduce difficulties in terms of drift and retaining high-fitness peaks even if/when discovered, besides severe penalties to average fitness). Importantly, these generation times are very high in comparison to the environmental switching frequencies that a learning population is able to thrive in.

The above model and experiments illustrate how one variety of phenotypic plasticity can enable a population to rapidly evolve across fitness valleys, a pattern of evolutionary change that cannot be experienced by non-plastic populations. On the flip-side, the benefits to the plastic population are lessened under more and more constant environments. At some point any benefits would be outweighed by the costs of learning (e.g., energetic cost of memory, risks). Although such aspects are omitted from the model here, and thus trade-offs are not directly visible in the results, the logic is straightforward: if genetic adaptation alone is sufficient in some stable environment, we should not expect to see plasticity playing any significant role.

Besides these two examples of environments, what other responses to changing environments might we expect to see in evolution? In the extreme case where many novel challenges appear within one lifetime, there may be plastic responses that do not become assimilated into the genome. In the absence of any regularity in those challenges, if the new challenges in one's lifetime are unrelated to the challenges faced by their ancestors, plasticity may be favourable but it is hard to see how any specific genetic adaptation would arise. If changes occur over a few generations, a Baldwinian-type interaction between plasticity and genetic evolution leading to genetic assimilation may result. Alternatively, if environmental changes are particularly repetitive, we may expect to see polyphenic/polymorphic genotypes and environmentally-sensitive switching (see West-Eberhard 1986, 2003) as mentioned above. If the environmental changes are strongly structured, we may additionally see modular architectures evolve in the genotypes (Parter et al. 2008) that are able to more quickly adapt to new challenges (Watson et al. 2014).

12.5 Conclusion

Our argument has been that focusing solely on entities, be they genes, species or ecosystems, is inherently problematic for conservation practice. We first argued that what we called *biodiversity fundamentalism* is untenable. It is both untenable as a thesis concerning the exclusive focus on one compositional hierarchy and as a thesis concerning the existence of a fundamental unit of conservation. Secondly, we argued that the genetic, taxonomic and ecological hierarchies for characterising biodiversity in terms of genetic organisation, taxonomic inclusiveness and ecological nestedness are, given their compositional ethos, insufficient to capture the functional dimension of biodiversity, particularly the evolutionary processes that maintain and originate new biodiversity units. Thirdly, we have proposed a *complementary entity-and-process-based approach to conservation practice*. Within this context, we distinguished between two types of important properties that evolutionarily significant units (i.e., ESUs) of conservation interest might exhibit: those amenable to conservation because they *preserve* natural dynamics (e.g., being a distinctive lineage) and those that *promote* them (e.g., being a population with a greater capacity for adaptation to change or stress). We focused on the latter because conservation strategies are aimed to identify not only “museums” but also “cradles” of biodiversity. Given this background, the hypothesis we wanted to test is whether plastic populations of a species might be considered ESUs with relevance for conservation. In particular, we wanted to show that plastic subpopulations that have enhanced evolutionary potential vis a vis non-plastic subpopulations make them amenable to ESU status. The model indeed shows that plasticity yields evolutionary potential, which is displayed in environments that switch in a few to a few tens of generations. Thus, populations with adaptation capacities available might possess an interesting property to consider when deciding on how to focus conservation efforts. Given that plastic populations might be important for species' adaptability and diversification,

they might be considered ESUs potentially amenable to conservation. This vindicates, on the one hand, a process-based approach to biodiversity and, on the other, suggests the need to take into account the processes generating plasticity when considering conservation efforts.

Acknowledgements Many thanks to Sandro Minelli, Philippe Huneman and Elena Casetta for feedback. Davide Vecchi acknowledges the financial support of the Fundação para a Ciência e a Tecnologia (Grant N. SFRH/BPD/99879/2014 and BIODECON R&D Project. Grant PTDC/IVC-HFC/1817/2014). Rob Mills acknowledges support by UID/MULTI/04046/2013 centre grant from FCT, Portugal (to BioISI).

References

- Angermeier, P. L., & Karr, J. R. (1994). Biological integrity versus biological diversity as policy directives: Protecting biotic resources. *Bioscience*, 44(10), 690–697.
- Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30, 441–451.
- Bradshaw, A. D. (1965). Evolutionary significance of phenotypic plasticity in plants. *Advances in Genetics*, 13, 115–155.
- Casetta, E., & Marques da Silva, J. (2015). Facing the big sixth: From prioritizing species to conserving biodiversity. In E. Serrelli & N. Gontier (Eds.), *Macroevolution: Explanation, interpretation and evidence* (pp. 377–403). Cham: Springer.
- Cook, S. A., & Johnson, M. P. (1968). Adaptation to heterogeneous environments. I. variation in heterophylly in *ranunculus flammula* l. *Evolution*, 22(3), 496–516.
- Dupré, J. (2012). *Processes of life: Essays in the philosophy of biology*. Oxford: Oxford University Press.
- Faith, D. P. (2016). Biodiversity. In *The Stanford encyclopedia of philosophy* (Summer 2016 edition), ed. Edward N. Zalta. <http://plato.stanford.edu/archives/sum2016/entries/biodiversity>. Accessed 25 Sept 2018.
- Fitzpatrick, B. M. (2012). Underappreciated consequences of phenotypic plasticity for ecological speciation. *International Journal of Ecology*. <https://doi.org/10.1155/2012/256017>.
- Forsman, A. (2015). Rethinking phenotypic plasticity and its consequences for individuals, populations and species. *Heredity*, 115, 276–284.
- Frankel, O. H. (1974). Genetic conservation: Our evolutionary responsibility. *Genetics*, 78, 53–65.
- Franklin, J. F. (1988). Structural and functional diversity in temperate forests. In E. O. Wilson (Ed.), *Biodiversity* (pp. 166–175). Washington, DC: National Academy Press.
- Grosberg, R. K., Vermeij, G. J., & Wainwright, P. C. (2012). Biodiversity in water and on land. *Current Biology*, 22(21), R900–R903.
- Hennig, W. (1966). *Phylogenetic systematics*. Urbana: University of Illinois Press.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1(3), 495–502.
- Luisi, P. (2003). Autopoiesis: A review and a reappraisal. *Naturwissenschaften*, 90, 49–59.
- Mace, G. M., Gittleman, J. L., & Purvis, A. (2003). Preserving the tree of life. *Science*, 300(5626), 1707–1709.
- Mayr, E. (1969). The biological meaning of species. *Biological Journal of the Linnean Society*, 1, 311–320. <https://doi.org/10.1111/j.1095-8312.1969.tb00123.x>.
- Mayr, E. (1970). *Populations, species, and evolution*. Cambridge: Harvard University Press.

- Mills, R., & Watson, R. A. (2005). Genetic assimilation and canalisation in the Baldwin effect. In M. S. Capcarrère et al. (Eds.), *Advances in artificial life* (pp. 353–362). Berlin: Springer.
- Mills, R., & Watson, R. A. (2006). On crossing fitness valleys with the Baldwin effect. In L. M. Rocha (Ed.), *Proceedings of the tenth international conference on the simulation and synthesis of living systems* (pp. 493–499). Cambridge: MIT Press.
- Miner, B. G., Sultan, S. E., Morgan, S. G., Padilla, D. K., & Relyea, R. A. (2005). Ecological consequences of phenotypic plasticity. *Trends in Ecology and Evolution*, 20(12), 685–692.
- Moritz, C. (1999). Conservation units and translocations: Strategies for conserving evolutionary processes. *Hereditas*, 130(3), 217–228. <https://doi.org/10.1111/j.1601-5223.1999.00217.x>.
- Morris, J. J., Lenski, R. E., & Zinser, E. R. (2012). The black queen hypothesis: Evolution of dependencies through adaptive gene loss. *MBio*, 3(2), e00036–e00012. <https://doi.org/10.1128/mBio.00036-12>.
- Norton, B. G. (2001). Conservation biology and environmental values: Can there be a universal earth ethic? In C. Potvin et al. (Eds.), *Protecting biological diversity: Roles and responsibilities*. Montreal: McGill-Queen's University Press.
- Noss, R. F. (1990). Indicators for monitoring biodiversity: A hierarchical approach. *Conservation Biology*, 4(4), 355–364.
- Osborn, H. F. (1896). Oytogenic and phylogenetic variation. *Science*, 4(100), 786–789. <https://doi.org/10.1126/science.4.100.786>
- Paenke, I., Kawecki, T. J., & Sendhoff, B. (2009). The influence of learning on evolution: A mathematical framework. *Artificial Life*, 15(2), 227–245.
- Parter, M., Kashtan, N., & Alon, U. (2008). Facilitated variation: How evolution learns from past environments to generalize to new environments. *PLoS Computational Biology*, 4(11), e1000206.
- Prado-Martinez, J., Sudmant, P. H., Kidd, J. M., Li, H., Kelley, J. L., Lorente-Galdos, B., Veeramah, K. R., Woerner, A. E., O'connor, T. D., Santpere, G., & Cagan, A. (2013). Great ape genetic diversity and population history. *Nature*, 499(7459), 471–475.
- Rosenberg, A. (1997). Reductionism redux: Computing the embryo. *Biology and Philosophy*, 12, 445–470.
- Ryder, O. A. (1986). Species conservation and systematics: The Dilemma of subspecies. *Trends in Ecology and Evolution*, 1, 9–10.
- Sarkar, S. (2002). Defining “biodiversity”; assessing biodiversity. *The Monist*, 85(1), 131–155.
- Sarkar, S. (2005). *Biodiversity and environmental philosophy: An introduction* (Cambridge studies in philosophy and biology). Cambridge: Cambridge University Press.
- Schaffer, J. (2003). Is there a fundamental level? *Noûs*, 37, 498–517.
- Simpson, G. G. (1953). The Baldwin effect. *Evolution*, 7(2), 110–117.
- Smith, T. B., Bruford, M. W., & Wayne, R. K. (1993). The preservation of process: The missing element of conservation programs. *Biodiversity Letters*, 1(6), 164–167.
- Staddon, J. E. R. (1983). *Adaptive behavior and learning*. Cambridge: Cambridge University Press. (Internet edition 2003).
- Sznajder, B., Sabelis, M. W., & Egas, M. (2012). How adaptive learning affects evolution: Reviewing theory on the Baldwin effect. *Evolutionary Biology*, 39(3), 301–310.
- Turney, P., Whitley, D., & Anderson, R. (1996). Evolution, learning, and instinct: 100 years of the Baldwin effect. *Evolutionary Computation*, 4(3), iv–viii. <https://doi.org/10.1162/evco.1996.4.3.iv>.
- Valladares, F., Matesanz, S., Guilhaumon, F., Araújo, M. B., Balaguer, L., Benito-Garazón, M., Cornwell, W., et al. (2014). The effects of phenotypic plasticity and local adaptation on forecasts of species range shifts under climate change. *Ecology Letters*, 17, 1351–1364. <https://doi.org/10.1111/ele.12348>.
- Waddington, C. H. (1953). Genetic assimilation of an acquired character. *Evolution*, 4, 118–126.
- Watson, R. A., Wagner, G. P., Pavlicev, M., Weinreich, D. M., & Mills, R. (2014). The evolution of phenotypic correlations and “developmental memory”. *Evolution*, 68(4), 1124–1138.

- West-Eberhard, M. J. (1986). Alternative adaptations, speciation, and phylogeny A review. *Proceedings of the National Academy of Sciences*, 83(5), 1388–1392.
- West-Eberhard, M. J. (2003). *Developmental plasticity and evolution*. Oxford: Oxford university press.
- Wimsatt, W. C. (2007). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Cambridge: Harvard University Press.
- Wolpert, L. (1994). Do we understand development? *Science*, 266, 571–572.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.