

Adaptationism and the adaptive landscape

Jon F. Wilkins · Peter Godfrey-Smith

Received: 28 October 2008 / Accepted: 29 October 2008 / Published online: 13 December 2008
© Springer Science+Business Media B.V. 2008

Abstract Debates over adaptationism can be clarified and partially resolved by careful consideration of the ‘grain’ at which evolutionary processes are described. The framework of ‘adaptive landscapes’ can be used to illustrate and facilitate this investigation. We argue that natural selection may have special status at an intermediate grain of analysis of evolutionary processes. The cases of sickle-cell disease and genomic imprinting are used as case studies.

Keywords Natural selection · Adaptationism · Adaptive landscape · Genetics · Optimality · Imprinting · Game theory

Introduction

Criticisms of the adaptationist perspective fall generally into two forms. First, critics have pointed out the relative ease of constructing an adaptationist hypothesis, contrasting that with the difficulty (or lack of interest) in constructing a rigorous test of that hypothesis. That is, critics argue that the standard of proof for an adaptationist explanation cannot simply be plausibility. This is a methodological concern, focusing on the standards of scientific argument.

The second criticism focuses on factors other than natural selection that contribute in important ways to evolution. One of these is the role of historical constraint, or path dependence in evolutionary processes. Features of organisms may not be optimal because they are derived from previous forms, which may have

J. F. Wilkins
Santa Fe Institute, Santa Fe, NM 87501, USA

P. Godfrey-Smith (✉)
Harvard University, Cambridge, MA 02138, USA
e-mail: pgs@fas.harvard.edu

evolved in a different context and for a different purpose. Critics also point to the importance of mechanistic factors. In many cases, the underlying genetic or developmental mechanisms may place inherent limitations on the degree to which an organism can adapt.¹

The debates over adaptationism are subsiding, with most researchers now taking a more nuanced view, acknowledging the centrality of adaptation as a concept in evolutionary analysis while at the same time recognizing the dangers inherent in pan-adaptationist reasoning. For example, many traits are currently studied in an explicit phylogenetic context. This approach considers the historical context in which a trait evolved. This framework is specifically geared towards separating out the effects of evolutionary change from those of shared history.

A second example is the field of population-genetics, where much of the past few decades have focused on the development and implementation of tests of neutrality (e.g., Tajima 1989; McDonald and Kreitman 1991; Fu and Li 1993). Implicit behind such a scientific program is a view that incorporates aspects of both the adaptationist and anti-adaptationist viewpoints. Neutral, or nearly neutral, evolution, dominated by genetic drift, is taken as a null hypothesis, and is thought to be sufficient to explain most of the genetic variation found in nature. In this sense, population-genetics grounds itself in an anti-adaptationist framework. However, against this background, the emphasis is decidedly on the identification of particular parts of the genome that show evidence of having recently been under selection. This emphasis implicitly treats the most interesting parts of evolution as those in which selection plays a strong role. The identification of selected regions of the genome is done in a statistical framework, which explicitly defines the standard of proof not as the ability to construct a plausible story about selection, but as the ability to statistically identify deviations from a well defined null model.

One area in which the disputes seem not to have subsided is evolutionary psychology, where strong adaptationist assumptions are still used and hotly debated (Barkow et al. 1992). But this is a field outside of evolutionary biology itself; the debate there concerns the best way to “import” well-established evolutionary ideas into a new theoretical context. This might also be seen as an area in which a simplistic view of evolution was initially assumed for legitimate Occamist purposes—how much can be explained with just one kind of evolutionary reasoning?—though this initial idealization has sometimes hardened into dogma.

While arguments about adaptationism no longer dominate the discourse in evolutionary theory itself, we believe that the issue has not been adequately understood and resolved. In particular, we believe that the best resolution of these debates is not just a vague middle ground, in which adaptation is viewed as an important, but not *the* important factor in evolution. While we would agree with this characterization, we believe that there are certain regularities that can help us to better understand the contexts in which an adaptationist perspective is particularly compelling and/or useful.

Here we offer a framework for rethinking part of the debate. The main idea we offer is the importance of *grain* in assessments of the role of selection in evolution.

¹ Gould and Lewontin’s classic (1979) paper used both of these lines of criticism.

On the one hand, we have coarser-grained perspectives on evolutionary processes. These perspectives consider longer time-scales and larger and more diverse sets of potential variants. On the other hand, more fine-grained perspectives focus on short-term evolutionary processes, considering a smaller set of potential variants in a higher degree of mechanistic detail. We argue that the role of selection in relation to other factors may be very different depending on which grain is considered. If there is a domain in which selection has a special role, it is at an *intermediate* grain of analysis.

In presenting these ideas we will use the conceptual framework of the “adaptive landscape.” The adaptive landscape has been widely used as a reasoning tool in evolutionary biology since it was introduced by Sewall Wright, but its use is also controversial. We discuss some of the limitations of the landscape metaphor below, as well as some of the caveats that must be heeded when this metaphor is employed. There are undoubtedly some readers whose skepticism regarding the adaptive landscape metaphor will incline them to dismiss any argument expressed in its terms. However, the main arguments presented here do not depend on the viability of the landscape idea, which is offered here simply as a vivid way of illustrating our framework.²

The adaptive landscape and its caveats

The metaphor of the adaptive landscape was introduced into evolutionary biology by Sewall Wright (1932). The landscape idea has always been controversial, and there exists a substantial literature detailing the various ways in which it is limiting or misleading.³ However, despite these well known limitations, we believe it to be a useful conceptual framework for considering the role of selection in relation to certain other factors in evolution.

In its simplest conception, the adaptive landscape represents a mapping from an entity’s characteristics to its fitness, with fitness represented in the vertical direction. To this may be added a second idea: natural selection can be thought of as a sort of hill-climbing algorithm, which results in movement by a population from areas of lower fitness to areas of higher fitness. Here we already encounter some of the complications associated with the landscape concept. What exactly are the entities occupying the landscape? How are the dimensions of the landscape defined, and what are the rules of movement? What determines the shape of the landscape? In the remainder of this section we will survey these problems, noting some of the pitfalls associated with using the landscape metaphor as a reasoning tool.

² The relationship between the present framework and some other aspects of the adaptationism debates—including simplicity preferences, model selection, and claims about which biological problems might be central or primary—is discussed in Godfrey-Smith and Wilkins (2008).

³ For discussions see Moran (1964), Gilchrist and Kingsolver (2001), Pigliucci and Kaplan (2006) and Calcott (2008).

Defining the entities occupying the landscape

In conceiving of an evolutionary problem in terms of a fitness landscape we must first specify the entities that we imagine occupying and moving on that landscape. These have variously been defined as organisms, populations, or species.

The “horizontal” axes of the landscape

The choice of entities constrains, but does not determine, the choice of the features represented on axes of the landscape. These may be individual-level properties (whether genetic or phenotypic) or population-level properties (such as the frequency of a gene). Even within the broad class of “phenotypic” traits, axes can be defined at a wide range of levels. An axis might represent the expression or activity level of a particular gene within a particular tissue. On the other hand, it could represent a macroscopic measurement, such as height, or probabilistic descriptions of behaviors.

Choice and number of dimensions

No one would deny that a complete description of any real biological system involves a ridiculously large number of dimensions. Invariably, researchers must focus on a small number of axes to render an analysis tractable. This amounts to an implicit assumption that any other aspects of phenotype are either fixed over the timescale of the analysis, or that any variation in other aspects of phenotype are largely irrelevant to the aspects of phenotype under consideration.

Neighborhood relations

The shape of a landscape depends on which possible states of organisms are close to which. This is relatively straightforward when the “horizontal” axes are genetic—either genotypes or gene frequencies. It is more contentious when we think of the axes as phenotypic, because “similar” phenotypes need not really be “close” to each other in the sense relevant to evolution; they need not be reachable by a relatively small number of mutational steps.

There are also some more specific dangers that must be kept in mind when using the adaptive landscape as a tool for reasoning. In each case, the problem is that the landscape metaphor implies certain characteristics that are at odds with the reality of biological systems.

Dimensionality

Because we inhabit a three-dimensional world, it is difficult for us to conceptualize or represent a landscape with more than two phenotypic dimensions. In fact, the term “landscape” implies an analogy with something like a topographical map, reinforcing this two-dimensional perspective. Any analysis will be limited in the number of dimensions that can be treated. However, the landscape metaphor and

representation may create a bias towards reducing the dimensionality of the problem below a level that is sensible for the questions being asked (Gavrilets 2004).

Invariability of fitness

The landscape metaphor implies that the fitness of a particular phenotype is the product of a fixed environment. This may be approximately true for a small subset of traits; for example, natural selection on enzymes involved in oxidative metabolism may be predominantly determined by the (approximately invariant) concentration of oxygen in the atmosphere. For most traits, however, natural selection is significantly, or even predominantly, determined by interactions with other organisms. In a sense, it may be useful to think not of organisms moving on a fixed landscape but on something more like a waterbed, where the height of the landscape at any point depends on the instantaneous distribution of organisms elsewhere on the landscape. A simple example of this type of effect is a trait with frequency-dependent fitness effects. The best-known example of a frequency-dependent fitness effect arises in the context of sickle-cell anemia, which we will discuss in detail below.

Further, in conflict with the “waterbed” metaphor, in biological systems these effects may be strongly non-local. For example, consider host-pathogen interactions, where the fitness landscape in the host region depends strongly on the distribution of organisms in the pathogen region. While the interaction between these two regions might be strong, there are probably few phenotypic dimensions along which the regions are not separated by extreme distances.

Niche construction

An additional complication to the point about landscape variability comes from the fact that many organisms modify their environments. Therefore, the shape of the fitness landscape depends not only on the instantaneous distribution of organisms on that landscape, but potentially on the history of how the landscape has been populated (Lewontin 1983; Odling-Smee et al. 2003).

Movement on the landscape

The idea of a landscape populated by hill-climbers implies that movement is local and continuous. While this may be a reasonable approximation in some cases, the appropriateness of this assumption depends on the choice of axes, the scale of analysis, and the nature of mutation, recombination, etc., within the particular system. In principle, any change in phenotype due to mutation or recombination will be discontinuous. At sufficiently coarse-grained scales of analysis, these discontinuities may be small enough to warrant a continuous approximation. On the other hand, at any scale of analysis there exists the possibility of macro-mutations with large phenotypic effects. Whether these mutations are sufficiently common to warrant explicit inclusion in an analysis will depend on the details of the system being analyzed.

A second issue raised by the idea of organisms being hill-climbers on a fitness landscape is that it may imply a type of intentionality to evolution. We believe that this is a greater danger in the context of popular representations of evolutionary processes than in technical analyses. It is also relatively easily avoided.

While we acknowledge the importance of explicitly keeping these caveats in mind, we do not believe that they constitute fatal flaws of the landscape metaphor. In the remainder of this paper we will use this metaphor to illustrate the role of grain in affecting the role of adaptation and other factors in evolution. In this context, we imagine a single landscape inhabited by many different kinds of organisms. We will focus on axes defined by phenotypic variables that describe individuals. This is in contrast with a landscape describing genetic properties of individuals or frequencies of genes in a population. However, the notion of “phenotype” used here is a broad one, including biochemical properties that might often be thought of as closer to the genotype.

When taken too literally, a use of the landscape metaphor like ours, in which a very diverse organisms are imagined on a single landscape, leads to incoherence. Below we will write as if the same measure of fitness can be applied to very different organisms. For analyses that focus on a limited region in the landscape, spanning, for instance, only a single species, this is not an issue. As we imagine zooming out to encompass a more and more diverse set of organisms, the problems of incommensurability get more and more acute. However, our aim is to make a point that does not depend on taking this spatial landscape metaphor fully seriously—our aim is to use the metaphor to illustrate features of the biology itself.

Three levels of grain

When one begins to consider an evolutionary problem, one must first choose a scale, or grain, of analysis. We can think of this as choosing how large a region of the adaptive landscape we want to include in the analysis. Choosing a larger region means considering a wider variety of alternatives (genotypes, phenotypes, strategies, etc.), but does not permit the same depth of analysis that could be performed on a smaller region. Also, if most evolutionary processes (e.g., mutation and selection) result in local movement on the landscape, then inclusion of a larger region implicitly considers these processes over a longer timescale.

To see the importance of grain in relation to attitudes towards adaptation, we will consider three different scales of analysis. At the highest, most “zoomed-out” level, evolutionary analyses consider broad patterns of occupancy across large regions of the landscape (see Fig. 1). At this level of analysis, populations are represented by single points, corresponding to typical features of the individuals in the population. (This creates problems in the case of extremely sexually dimorphic species, like some barnacles, but we will idealize away from that problem.) At this level of grain, natural questions to ask include the following: how are actual organisms distributed in the space of possible organisms? What proportion of the landscape’s peaks are occupied? To what extent are the occupied peaks the highest peaks? How thoroughly and how predictably does natural selection explore the landscape?




Fig. 1 An adaptive landscape from the most coarse-grained perspective. In each of these figures, the *horizontal axes* are phenotypic variables and the *vertical* is fitness. *Points* indicate locations of extant species or populations. While these points are found at or near adaptive peaks, they are not necessarily the highest peaks. The occupation of the landscape is sparse, indicating that a large number of conceivable phenotypes may never have been explored. The ruggedness of the landscape suggests that it may be impossible for a population to move from one peak to a neighboring peak, even if that neighboring peak is substantially higher




Fig. 2 A view of the adaptive landscape at an intermediate level of grain. Here we focus on an area of the landscape from Fig. 1 that contains a single peak. By restricting consideration to this narrow region, neighboring peaks and large unoccupied regions have been excluded. From this perspective, there appears to be a single optimal solution, and the population, which is sitting at the peak, appears to have found that solution. Note that as we have moved closer, the point representing the population has become more diffuse

To most biologists working at this highest level of analysis, what is most striking is the emptiness of the landscape. In the history of life on earth, organisms have explored a vanishingly small fraction of the conceivable ways of making a living, and the idea that the modes of life that we see today somehow represent the “best” of these possible forms has been broadly rejected. Rather, populations have been restricted to a small subset of local peaks by chance, as well as by historical and developmental constraints. From this vantage point, the power of natural selection to produce adaptation appears quite limited.

The situation changes if we zoom in on a particular region of the landscape, perhaps containing only one or a few peaks. At this intermediate level of grain,

whole populations still tend to occupy single points, but the points are vague or smudged (see Fig. 2). These analyses also typically focus on variation in a small number of dimensions, assuming (often implicitly) that the traits represented by the other dimensions are invariant over the timescale of the analysis, and evolve independently of the trait or traits under consideration.

A question asked in this second context might take the form, “Given that there is a population somewhere in this region of the landscape, where, more exactly, should we expect to find it?” The answer that many biologists would give is that we expect to find the population at or near one of the local peaks. Given enough time and a local topography that is conducive to a thorough exploration of the region, we might even expect to find the population at or near the highest of those local peaks. Perhaps more commonly, however, research at this scale starts with the empirical observation that a population occupies a particular location within the region. The task is then to uncover whether and why there should be a peak in that location.

At this intermediate scale, many will hold that the salient feature of evolution is local adaptation. Populations tend to be found near peaks, as opposed to in the adjacent valleys. Some biologists may disagree with that claim, but the crucial point is that selectionist conclusions drawn at this level of analysis do not contradict non-selectionist conclusions drawn at the higher level. The fact that the population is in a particular general region of the landscape, as opposed to some very distant one, is simply one of the background assumptions made when working at this second level of analysis.

Now let us consider a third level of “zoom” with which we might view the landscape. We now focus in great detail on a very small region. When we do this, certain features of evolution that are typically ignored at the higher levels become critically important. We are now looking so closely that a population does not inhabit a single location, or a compact area, in the landscape. Instead, we have reached a grain of analysis in which individuals themselves can be recognized (see Fig. 3). These individuals will have a distribution on the landscape, and the distribution is subject to complex processes of change. Analysis of the evolutionary dynamics of the system must account for drift and mutation, as well as selection, and, in the case of diploid organisms, Mendelian inheritance and recombination. The way the population moves on the landscape from generation to generation is not continuous, at this level of grain. Points appear suddenly, some distance from their parents.⁴

As we focus more on the details of evolution at this lowest level, the adaptationist features that were prominent at the intermediate level of analysis recede. Evolutionary processes involving random fluctuations, interactions among alleles or loci, and constraints inherent in the mechanism of inheritance become important aspects of our understanding. Here again, there will be some disagreement among

⁴ In one sense, molecular work is the “finest-grained” work in evolutionary biology, because it works at the scale of molecules and cells. But this is not the same of “grain” as the sense we employ here. Depending on assumptions made about time-scale and the range of variants considered, molecular evolutionary work can be at either the finest or intermediate levels in our sense. Molecular evolution, as noted in our “Introduction”, is a part of biology that has successfully integrated adaptationist and non-adaptationist approaches.




Fig. 3 A fine-grained view of the landscape. We have zoomed in to focus on the region of the peak from Fig. 2 that is occupied by the population. In contrast with Figs. 1 and 2, where a single mark was used to indicate the population as a whole, we now use *points* to indicate the locations of individual organisms. Here a different set of limitations on adaptation become salient, including discontinuous movement on the landscape, drift, recombination, and so forth

biologists about which factors are in fact most significant. But many will hold that at this lowest level, as we also saw at the highest level, what is often most salient is the set of features of evolution that frustrate adaptation. And once again, the crucial point is that de-emphasizing selection at this lowest level is not at all inconsistent with applying a strongly selectionist approach at the second, middle level.⁵

The sickle-cell case

We will illustrate these ideas with two examples. The first is very familiar. Sickle-cell disease afflicts individuals who inherit two copies of a mutant allele of the Hemoglobin alpha chain gene. Despite the devastating effects of this condition, this mutant allele is present at an appreciable frequency in certain populations, because heterozygous individuals, who carry one copy of the mutant (HbS) allele and one copy of the normal (HbA) allele, have an increased resistance to malaria. It is instructive to consider how different scales of analysis of this system drive different relationships to adaptationism.

We can first consider an analysis at the finest grain, corresponding to the single-generation timescale. In an environment where malaria is present, there exists an optimal genotype, the HbA/HbS heterozygote, which protects against malaria without other significant negative consequences. However, a population composed entirely of “adapted” heterozygotes is inherently unstable (see Fig. 4). Mendelian inheritance reconstitutes both the HbA/HbA heterozygote, which is susceptible to

⁵ The distinction between evolution on shorter time-scales in which the set of available variants is fixed and longer time-scales in which the set of variants is not fixed has been treated dynamically by Eshel and Feldman (1984, 2001), Hammerstein (1996), and Nowak and Sigmund (2004), with conclusions that complement the present analysis. In effect, this is a tradition of explicit modeling of the relations between the finest and intermediate grains of analysis considered here. These models often, in an understandable idealization, treat the processes at the intermediate level of grain as sequential combinations of processes at the finer level.

Fig. 4 Fine-grained view of the sickle-cell case. On the timescale of a single-generation, we find that the well adapted HbA/HbS heterozygote is disrupted each generation, and Mendelian segregation reconstitutes both homozygotes





Fig. 5 The sickle-cell case viewed at an intermediate timescale. Here we imagine a hypothetical error in recombination at the hemoglobin α -chain locus that creates a new HbAS allele that encodes both the HbA and HbS polypeptides. Such an allele might provide protection against malaria without causing sickle-cell disease, even when homozygous

malaria, as well as the HbS/HbS heterozygote, which results in sickle-cell disease. When we take this fine-grained perspective, we find that our system of genetic inheritance continually undermines humans' ability to adapt to the presence of malaria. This sort of phenomenon is common when we look at evolution in a fine-grained way; mechanistic considerations often limit the ability of natural selection to produce adaptation.

At a slightly coarser grain of analysis, we begin to see signs of adaptation. If we expand our view to encompass the mutation event that created the mutant allele, we see that human evolution has found the beginnings of a solution to the problem of malaria. From an adaptationist point of view, heterozygous resistance to the disease certainly represents an advance over a population where everyone is susceptible. If we expand, hypothetically, to an even longer timescale, we can imagine solutions to the problem that are not disrupted by the diploid genetic architecture. For example, the right mistake in recombination could create a chromosome that carried both the normal and mutant alleles (Fig. 5). This new allele might confer malaria resistance to its carriers without producing sickle-cell anemia in homozygotes.

As it happens, natural populations in Africa do, in fact, contain an allele (HbC) that confers resistance in homozygous state without harmful sickling, but this allele has low fitness in heterozygous state with each of the prevalent alleles (HbA and HbS), so it has been unable to advance from very low frequency (Tempelton 1982; Gilchrist and Kingsolver 2001). There we see the role of a different kind of fine-

grained constraint, the requirement that a new favorable allele be advantageous when appearing as a heterozygote with locally common alleles. Again, when operating within a coarser-grained perspective we can expect new alleles to eventually arise that overcome this constraint.

If we continue to zoom out, the appearance of adaptation begins to recede again. Why, for instance, do we not have a fundamentally different immune system that would make malaria a non-issue? We can certainly conceive of a species very much like us, but with a differently structured immune system that would defeat malaria. Our susceptibility to malaria suggests a whole other class of features limiting our adaptation; we are subject to numerous historical and developmental constraints. Our immune system, for example, is based on a structure that has been in place for hundreds of millions of years. Some features of the malaria case encountered when we take this most zoomed-out perspective also recall one of our earlier caveats about the adaptive landscape itself. On this longer timescale, the topography of the landscape is certainly not stable, and a key component of the environment—the malaria parasite—is itself evolving. If our immune systems were based on a fundamentally different structure, a host of pathogens, including malaria, would have been under selection to develop mechanisms to evade that system. In general, our capacity to adapt to our environment is limited by the fact that many aspects of our environment are simultaneously adapting to us.

Genomic imprinting

As a second example, we consider the phenomenon of genomic imprinting, and two different mathematical approaches to modeling the predominant evolutionary explanation for its origin, the kinship theory of imprinting. Briefly, genomic imprinting refers to the phenomenon where the pattern of expression of an allele depends on its parent of origin (for recent reviews, see Hore et al. 2007; Sha 2008). For instance, alleles at an imprinted locus might be expressed when maternally derived, but transcriptionally silent when paternally derived. The kinship theory proposes that imprinting results from asymmetries in the way natural selection is acting on maternally and paternally derived alleles (Haig 2002; Wilkins and Haig 2003). The expression strategy that maximizes the inclusive fitness of an allele can vary depending on the allele's parent of origin.

Many imprinted genes are expressed in fetal or placental tissues, and are associated with growth effects and modulating the demand for maternal resources. In these cases, it is easy to understand the basis of the inclusive fitness asymmetry. Increased growth and resource demand presumably increase individual fitness, but induce a cost on the mother's other offspring. A maternally derived allele will be present in half of those other offspring, so its optimal expression strategy will balance these costs and benefits in a way that maximizes the transmission of that allele to future generations. It may be prudent, from this point of view, to withhold resources from one offspring in order to retain the mother's good condition for future reproduction. The mother's other offspring may or may not share the same

father, however, so a *paternally* derived allele in a present fetus should place a lower weight on those indirect costs. Therefore, an allele will favor placing a greater demand on maternal resources when it is paternally derived than when it is maternally derived.

Most of the mathematical models of the evolution of imprinting fall into two classes. The first class of models relies on formal population-genetic methods that explicitly consider the selective pressures acting on two different alleles segregating in the population (e.g., Spencer et al. 1998). A two-allele model with imprinting includes four distinct genotypes. Pairs of offspring from the same mother can have sixteen different combined genotypes. These models parameterize the fitness effects associated with each of these possible situations, and then consider the effect of natural selection on the frequencies of alleles. The result is a set of relationships among parameter values that define when a novel allele can invade the population, and when it can be expected to sweep to fixation.

The second class of model relies on stability analysis. Basically, the goal is to find an evolutionarily stable strategy (ESS). This is a strategy which, once established, is resistant to invasion by alternatives (Maynard Smith 1982). Here, the expression level of an allele is treated as a continuously variable quantity. The ESS is defined as a set of maternal and paternal strategies with the property that neither allele can increase its inclusive fitness by adopting an alternative strategy.⁶ The basic results are that, for a locus where increasing expression results in increased fetal growth (and increased demand on maternal resources), the ESS has the maternally derived allele transcriptionally silenced, and the paternally derived allele expressed at the level that maximizes its inclusive fitness. Conversely, the paternally derived allele will be silenced at a locus where increased expression suppresses growth, and reduces demand on the mother (Wilkins and Haig 2001).

A significant debate has arisen over the fact that the two types of analysis yield apparently contradictory conclusions. ESS-based stability analysis predicts that an imprinted growth enhancer will always be maternally silenced, and that the only condition required for natural selection to favor this solution is a non-zero probability of multiple paternity. The population-genetic analysis, on the other hand, clearly identifies sets of parameter values for which paternal silencing of a growth enhancer can invade the population. Furthermore, it is possible for natural selection to favor an imprinted allele over an unimprinted one even in the absence of multiple paternity.

The key to reconciling these contradictory results resides in the fact that the two modes of analysis are implicitly considering different grains of analysis, and different timescales over which natural selection is assumed to work. Specifically, the population-genetic analysis focuses on the fate of a particular imprinted allele when it is introduced into a population that is fixed for an unimprinted allele (one that has the same expression strategy irrespective of its parental origin). The ESS analysis, by contrast, works on a coarser grain of analysis, implicitly assuming that

⁶ Much population genetic modeling is concerned with finding states of populations that are stable (in the short-term sense in which only a fixed range of alleles are considered). An interest in stability is common to both approaches. But the population genetic approach can also consider the details of the dynamics, the role of stochastic effects, and the like, whereas the ESS approach *only* identifies stable states.

many alternative alleles have been introduced into the population over a long period of time. The population-genetics approach permits a detailed description of the dynamics of a system with particular alleles present. The ESS approach attempts to identify strategies that outcompete any alternative strategy (within a prescribed set). What is lost in the ESS approach is the ability to reconstruct how and whether that strategy could have come into existence and reached an appreciable frequency.

The superiority of one approach over the other will depend, then, on the nature of the questions being asked. The population-genetics approach gives an exact model of the dynamics of particular alleles. The ESS approach identifies the overall profile of a strategy that should be able to outcompete (once established) a specific range of alternatives. A researcher more interested in short-term behavior will choose the finer-grained analysis. The choice will also depend on whether the goal is to understand the dynamics of the evolutionary process, or just its outcome.

The ways in which different assumptions about time-scale affect the value of particular methods can be quite subtle. We will illustrate this by looking in more detail at the apparent contradiction between ESS and population-genetics approaches in the imprinting case. As noted earlier, there are circumstances in which the population-genetic model predicts that a paternally silenced imprinted allele at a growth enhancer can invade a population in which an unimprinted allele is at fixation. But what circumstances are these? It appears that this case arises when the expression level of the unimprinted allele is greater than what is optimal for either the maternally or paternally derived allele. In this case, any allele that reduced overall expression from the locus would be favored by natural selection, including an allele that is transcriptionally silent when paternally derived. That is, if paternal silencing is selectively favored, an unimprinted allele with reduced expression would also be selectively favored.

This reveals a hidden assumption behind the ESS approach. These analyses implicitly assume that the unimprinted state that the system starts in is one that corresponds to the optimal expression level for an unimprinted allele, which is less than the optimum for a paternally derived allele. If we imagine a situation in which the unimprinted expression that is greater than either optimum, the ESS analysis of imprinting assumes that natural selection will reduce the expression level below the paternal optimum *before* the introduction of an imprinted allele.

So in effect, the ESS approach assumes that a mutation that quantitatively changes the expression level of an allele is orders of magnitude more likely than a mutation that gives rise to imprinted gene expression. It also assumes that natural selection is sufficiently strong (relative to drift) that this expression level is effectively optimized before the introduction of the first imprinted allele. This particular set of assumptions seems reasonable, suggesting that the conclusions arising from the ESS analysis are likely to be valid. However, in the absence of the apparently contradictory population-genetic results, it would have been difficult to identify what the relevant hidden assumptions are, or to understand the consequences for the evolution of the system if those assumptions were violated.

Analysis and conclusions

Let us note explicitly some of the features of the three levels of grain that make adaptive features of evolution more, or less, prominent. At the finest level of grain, existing features of the genetic architecture are treated as fixed constraints against which selection must act. At the intermediate level, we suppose that some of these can be altered by such things as modifier alleles. At the intermediate level, time-scales are also longer than they are at the finer level, and we can suppose that there is a constant steady flow of new variants arising in the population. So the likely fate of any *particular* mutation is not important.

As we move across multiple levels of grain, we consistently see two effects. First, coarse-graining the analysis (increasing the timespan) increases the opportunities for natural selection to perform its evolutionary search. However, the same coarse-graining also increases the space that must be searched. It seems that at the finer levels of grain, the size of the search space is not limiting, and coarse-graining results primarily in an enhanced efficacy of search, corresponding to an apparent predominance of adaptation. At some point, that balance shifts, and the expansion of the search space outpaces the ability of selection and mutation to perform its search. We are then left with a strong sense of the path-dependencies and historical constraints that have shaped the diversity of life.

The way we have described the distinction between the intermediate and coarse-grained levels can also be thought of as a form of “ascertainment bias.” If we consider starting from the highest level and zooming in to the intermediate level, it is obvious that we do not do so at random. We usually focus on a region that is populated by an observed species. To the extent to which natural selection is effective at local optimization, this will tend to be a locally fit peak surrounded by less fit valleys. However, if we recall one of the caveats raised about the adaptive landscape metaphor, we see that there is another, perhaps more important, aspect of modeling bias that determines the effective grain of analysis.

Adaptive landscapes are typically conceptualized and modeled in a small number of dimensions, a number of dimensions likely to be orders of magnitude less than the number that would be required to provide a full description of an organism’s phenotype. In constructing a model, biologists choose which of these many dimensions to include. Phenotypic traits differ in their evolutionary lability, and the same timescale may correspond to different grains of analysis for different aspects of phenotype. It is likely the case that individual researchers tend to focus on those aspects of phenotype that are most amenable to the tools with which they are most comfortable. For instance, a behavioral ecologist looking at a system may be inclined to focus on those aspects of phenotype for which it is possible to quantitatively describe a single, smooth peak in the fitness landscape. To take a famous example, he might consider the relationship between beak size and the availability of different types of seeds. A population-geneticist looking at the same system on the same timescale might focus instead on non-coding sections of the genome, where she could employ all of the mathematical formalism that has been developed for analyzing populations evolving under neutrality.

Our suggestion, then, is that many questions about the adaptive character of evolution depend crucially on the grain at which evolutionary processes are being described. One way to represent various kinds of adaptationist commitment is as claims about the size and location of the region, on a scale representing these different levels of grain, within which evolutionary change tends to have an adaptive character. Some will say the region is large, others will insist that it is small. But we think that without this conceptualization of the situation, or something akin to it, it is difficult to frame the debate in a way that makes it tractable.

We offer this framework as a way of visualizing some of the preconceptions and assumptions that are common within different subcommunities of evolutionary biology, and the way that the use of different tools tends to highlight different aspects of the evolutionary process. We do not mean to imply that real adaptive landscapes tend to look like the ones we have described here, nor that the landscape metaphor is, in general, the most appropriate and useful way of thinking about evolution. It might be better to think of the adaptive landscape as described here as a model of the models that evolutionary biologists carry around in their brains. As in all cases of modeling, the tools used in our analyses emphasize some aspects of the process being modeled and deemphasize others. However, we believe that this framework may be particularly useful in reconciling apparently contradictory interpretations of biological systems.

Acknowledgment We are grateful to Patrick Forber for helpful comments on an earlier draft.

References

- Barkow J, Cosmides L, Tooby J (eds) (1992) *The adapted mind: evolutionary psychology and the generation of culture*. Oxford University Press, New York
- Calcott B (2008) Evaluating the fitness landscape revolution. *Biol Philos* (forthcoming)
- Eshel I, Feldman M (1984) Initial increase of new mutants and some continuity properties of ESS in two-locus systems. *Am Nat* 124:631–640. doi:[10.1086/284303](https://doi.org/10.1086/284303)
- Eshel I, Feldman M (2001) Optimality and evolutionary stability under short-term and long-term selection. In: Orzack S, Sober E (eds) *Adaptationism and optimality*. Cambridge University Press, Cambridge, pp 161–190
- Fu XY, Li WH (1993) Statistical tests of neutrality of mutations. *Genetics* 133:693–709
- Gavrilets S (2004) *Fitness landscapes and the origin of species*. Princeton University Press, Princeton
- Gilchrist G, Kingsolver J (2001) Is optimality over the hill? The fitness landscapes of idealized organisms. In: Orzack S, Sober E (eds) *Adaptationism and optimality*. Cambridge University Press, Cambridge, pp 219–241
- Godfrey-Smith P, Wilkins J (2008) Adaptationism. In: Sarkar S, Plutynski A (eds) *A companion to the philosophy of biology*. Blackwell, London, pp 186–201
- Gould SJ, Lewontin RC (1979) The spandrels of San Marco and the panglossian paradigm: a critique of the adaptationist program. *Proc R Soc Lond* 205: 581–598. Reprinted in Sober (1994), pp 73–90
- Haig D (2002) *Genomic imprinting and kinship*. Rutgers University Press, Piscataway
- Hammerstein P (1996) Darwinian adaptation, population genetics and the streetcar theory of evolution. *J Math Biol* 34:511–532. doi:[10.1007/BF02409748](https://doi.org/10.1007/BF02409748)
- Hore TA, Rapkins RW, Graves JAM (2007) Construction and evolution of imprinted loci in mammals. *Trends Genet* 23:440–448. doi:[10.1016/j.tig.2007.07.003](https://doi.org/10.1016/j.tig.2007.07.003)
- Lewontin RC (1983) *The organism as the subject and object of evolution*. Reprinted in Levins R, Lewontin RC, *The dialectical biologist*. Harvard University Press, Cambridge, 1985, pp 85–106

- Maynard Smith J (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge
- McDonald JH, Kreitman M (1991) Adaptive protein evolution at the ADH locus in *Drosophila*. *Nature* 351:652–654. doi:[10.1038/351652a0](https://doi.org/10.1038/351652a0)
- Moran PAP (1964) On the non-existence of adaptive topographies. *Ann Hum Genet* 27:383–393
- Nowak M, Sigmund K (2004) Evolutionary dynamics of biological games. *Science* 303:793–799. doi:[10.1126/science.1093411](https://doi.org/10.1126/science.1093411)
- Odling-Smee J, Laland K, Feldman M (2003) *Niche construction: the neglected process in evolution*. Princeton University Press, Princeton
- Orzack S, Sober E (2001) *Adaptationism and optimality*. Cambridge University Press, Cambridge
- Pigliucci M, Kaplan J (2006) *Making sense of evolution: the conceptual foundations of evolutionary biology*. University of Chicago Press, Chicago
- Sha K (2008) A mechanistic view of genomic imprinting. *Annu Rev Genom Hum Genet* 9:197–216. doi:[10.1146/annurev.genom.122007.110031](https://doi.org/10.1146/annurev.genom.122007.110031)
- Spencer HG, Feldman MW, Clark AG (1998) Genetic conflicts, multiple paternity and the evolution of genomic imprinting. *Genetics* 148:893–904
- Tajima F (1989) Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics* 123:585–595
- Tempelton AR (1982) Adaptation and the integration of evolutionary forces. In: Milkman R (ed) *Perspectives on evolution*. Sinauer, Sunderland, pp 15–31
- Wilkins JF, Haig D (2001) Genomic imprinting of two antagonistic loci. *Proc R Soc Lond B Biol Sci* 268:1861–1867. doi:[10.1098/rspb.2001.1651](https://doi.org/10.1098/rspb.2001.1651)
- Wilkins JF, Haig D (2003) What good is genomic imprinting: the function of parent-specific gene expression. *Nat Rev Genet* 4:359–368. doi:[10.1038/nrg1062](https://doi.org/10.1038/nrg1062)
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding and selection in evolution. In: *Proceedings of the sixth international congress of genetics*, vol 1, pp 356–366