



Contents lists available at ScienceDirect

Studies in History and Philosophy of Biological and Biomedical Sciences

journal homepage: www.elsevier.com/locate/shpsc

Confirmation and explaining how possible

Patrick Forber

Philosophy Department, Tufts University, Miner Hall, Medford, MA 02155, USA

ARTICLE INFO

Article history:

Received 19 June 2008

Received in revised form 24 September 2009

Available online xxxxx

Keywords:

Confirmation

Explanation

Evolutionary biology

Modeling

Possibility

ABSTRACT

Confirmation in evolutionary biology depends on what biologists take to be the genuine rivals. Investigating what constrains the scope of *biological possibility* provides part of the story: explaining how possible helps determine what counts as a genuine rival and thus informs confirmation. To clarify the criteria for genuine rivalry I distinguish between *global* and *local* constraints on biological possibility, and offer an account of how-possibly explanation. To sharpen the connection between confirmation and explaining how possible I discuss the view that formal inquiry can provide a kind of confirmation-theoretic support for evolutionary models, and offer an example of how-possibly explanation interacting with testing practice.

© 2010 Elsevier Ltd. All rights reserved.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Biological and Biomedical Sciences*

1. Introduction

The climate fluctuates and the availability of food resources changes. The beak of a Galapagos finch species evolves. A correlation between the environmental change and the evolutionary change can be evidence for natural selection. The data count as evidence by supporting the natural selection hypothesis over other rivals, such as hypotheses that invoke drift or constraint. If we take seriously this contrastive nature of confirmation, as Salmon (1990) and Sober (1990) argue we should, then evidential relations depend on a contrast class or testing set. Yet what rival hypotheses should biologists include in (or exclude from) the testing set? I aim to investigate what guides this kind of decision by analyzing minimal explanatory conditions that constrain the scope of *biological possibility*.

To gain traction on the problem, I will use the distinction between *how-possibly* and *how-actually* explanations, introduced by Dray (1957) and applied to evolutionary biology primarily by Brandon (1990). Prima facie, how-possibly explanations provide a guide to what counts as biologically possible. Yet there is a crucial ambiguity regarding the sorts of constraints introduced by explaining how possible (Sect. 2). There are broad *global* constraints, informed by formal inquiry into models of evolution, about what sorts of

evolutionary processes may occur and the potential patterns these processes can produce. There are also narrower *local* constraints, informed by empirical inquiry into real biological systems, about whether evolutionary processes can produce specific outcomes, such as the camera eye, wings, or a change in finch beak size. I shall offer an account of how-possibly explanation in evolutionary biology that respects the difference between global and local constraints, and helps vindicate the practice of constraining the focus of inquiry to relatively few rivals (Sect. 3). I then connect my account to previous discussions on how-possibly explanation (Sect. 4). Finally, I make two points in support of the connection between confirmation and explaining how possible. First, distinguishing global from local constraints clarifies the claim that formal analyses of abstract evolutionary models provide a kind of confirmation-theoretic support for those models (Sect. 5). Second, some historical controversies are usefully interpreted as expanding the scope of biological possibility by defending novel global how-possibly explanations. The origin of the neutral theory of molecular evolution provides an illustrative example of this (Sect. 6). How-possibly explanations thus inform confirmation by showing that some hypotheses, whether general proposals about evolutionary processes or specific models of some target system, meet the minimal conditions to be considered a genuine rival.

E-mail address: patrick.forber@tufts.edu

1369-8486/\$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.shpsc.2009.12.006

Please cite this article in press as: Forber, P. Confirmation and explaining how possible. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2010), doi:10.1016/j.shpsc.2009.12.006

2. Global versus local

One way to understand what counts as a biological possibility appeals to the distinction between *how-possibly* and *how-actually* explanations, due originally to Dray (1957). This is a familiar distinction, but it turns out to hide an ambiguity. So I will first lay out the standard view, then locate the ambiguity in question.

How-actually explanations aim to explain how or why some event actually occurs; these are the sorts of explanations usually considered in the explanation literature (Salmon, 1989; Woodward, 2009). In the context of evolutionary biology a how-actually explanation picks out the biological possibility that, given suitable evidence, we take to explain the target outcome or pattern. Kettlewell (1955, 1956) gives a how-actually explanation for the evolution of melanism in *Biston betularia*—the melanic phenotype confers a greater degree of camouflage that helps individuals avoid predators better than the non-melanic phenotype.¹ Brandon (1990) discusses the evolution of heavy metal tolerance in *Arabidopsis* to facilitate their growth on toxic mine tailings as another good example of a how-actually explanation in evolutionary biology. Both examples provide an explanation for how evolution unfolded in biological populations to produce the target traits, melanism and heavy metal tolerance. Also, these explanations are ostensibly well confirmed hypotheses about how natural selection *actually* produced specific evolutionary outcomes in real populations of organisms.

In contrast, how-possibly explanations aim to explain how some event could possibly occur. As Dray puts it, the function of a how-possibly explanation 'is to rebut the presumption that what happened was impossible, or at any rate extremely unlikely given the circumstances' (Dray, 1993, p. 27). This formulation deserves some clarification. First, a how-actually explanation would, of course, rebut such a presumption as well. A how-possibly explanation does so without providing the complete actual explanation. Second, Dray complicates the picture by adding that how-possibly explanations can show that some event is not 'extremely unlikely'. An unlikely event is still possible. Dray may be closer to scientific practice since scientists often exclude the very improbable from the set of genuine rivals. However, if *contingency* characterizes much of the evolutionary process, as Gould (1989) and Beatty (1995, 1997) argue, then such improbable events may play an important role and so should not be neglected. To be precise, we should correct Dray and count 'extremely unlikely' events as possible. The relative importance of contingency in the evolutionary process is an open question and interacts with explanation in a different way. It affects the *counterfactual resiliency* of our how-actually explanations for evolutionary phenomena, and thus whether we can give a *robust process* or *actual sequence* explanation of some target phenomenon (Jackson & Pettit, 1992; Sterelny, 1996). Counterfactual resiliency simply does not apply to how-possibly explanations, for an event is possible if it occurs in one possible state of affairs or many. Subtleties aside, Dray's notion of explaining how possible plays a vital role in evolutionary biology, as many have argued (Lewontin, 1985, 2000; O'Hara, 1988; Resnik, 1991; Plutynski, 2004, 2005).

Before moving on, let me set aside some issues about possibility and explanation. What is biologically possible depends on the way the world is, on what physical or biological laws hold, if any. If one takes Van Fraassen's (1977) perspective then biological possibility is a subset of logical or verbal possibility constrained by the biolog-

ical and physical laws. Yet the metaphysical nature of possibility, the proper account of laws in science, and even the existence of biological laws are all controversial issues. I will set these problems aside in order to investigate how biologists both develop and deploy their background theoretical framework to constrain what counts as a biological possibility. What really matters is what practitioners *take to be* possible given their background theoretical framework. Levi (1988) provides a clear way to approach this: explaining how possible involves determining whether some event is possible *relative* to a suitably constrained set of background information. While what is biologically possible is determined by the causal structure of the world, what biologists take to be possible depends on their accepted theoretical framework, and this can change as the framework changes (see Sect. 6). The overall theory of explanation for biology, and whether how-possibly explanations count as genuine explanations on this theory, are not resolved either.² I will take a pluralist perspective on explanation in science and, based on arguments made by Lewontin and others, treat how-possibly explanations as worthy of philosophical investigation. They help us answer a crucial question: *what are the biological possibilities?*

There are, however, *two* distinct ways to take this question about biological possibility. The first way takes the question to address the nature of evolutionary processes, focusing on the power and limits of selection, drift, and constraint operating in ideal populations. Can weakly adaptive traits evolve by natural selection in small populations? Can random drift produce a constant rate of molecular change in DNA and proteins? This sort of question concerns the scope of abstract, *global* biological possibility space, for it investigates the capabilities of general evolutionary models. The second way takes the question to address whether evolution can produce target outcomes in real populations. Dawkins (1986) and Nilsson and Pelger (1994) offer explanations of this kind for how the vertebrate camera eye may have evolved. Can natural selection for coping with informationally demanding foraging problems explain the evolution of human cognition? Can constraint explain the pattern of allometry observed across biological lineages? Can random drift explain the rate of protein evolution in Cytochrome C? This sort of question concerns the scope of concrete, *local* biological possibility space, for it guides speculation on how an evolutionary process can produce target traits and patterns.

The global question uncovers a formal kind of constraint, whereas the local question uncovers an applied kind of constraint. Abstract or mathematical investigation into how process models of evolution operate in *idealized* populations to produce potential evolutionary outcomes constrains the scope of global possibility, whereas concrete or empirical speculation on how those evolutionary processes can operate in *real* populations to produce specific target outcomes constrains the scope of local possibility. The mathematical or formal structure of models determines the global possibilities. In addition to the models, the background biology of a target population determines the local possibilities for that system. Local possibility space is a subset of global possibility space—global possibilities must be consistent with information about the accepted set of formal evolutionary models, whereas local possibilities must be consistent with that information plus information about some real biological system. Call the first body of information the *global information set*. For simplicity I will assume that there is one global information set shared across the field. Call the second body of information, enriched by specific biological de-

¹ There is some debate about the results of Kettlewell's studies (Hagen, 1999; Rudge, 1999). Nevertheless, if the studies were to provide the results claimed using the proper control groups then these studies would give a how-actually natural selection explanation.

² Reiner (1993), for example, argues that how-possibly explanations should not count as true explanations because they merely resolve puzzlement and are incomplete.

tails of a target system, the *local information set*. A local information set is determined relative to a target biological system.³

Concerns about the biological world guide investigation into the scope of both global and local possibility spaces. Yet such concerns guide the abstract formal modeling in a way very different than the concrete application. Formal inquiry focuses on idealized populations of abstract entities that do not make contact with any particular biological lineage. Model building requires idealization, and while empirical concerns may guide or support these idealizations, they need not. In fact, models with no empirical connections or support often play an important role, helping us understand constraints on related processes (Wimsatt, 2002). On the other hand, concrete application brings evolutionary models in contact with particular biological lineages. What is possible in the general evolutionary sense may not be possible for a particular lineage; incorporating specific empirical details from a target system introduces further constraints on what we take to be a biological possibility for that lineage. The distinction between global and local questions about biological possibility accounts for this difference.

3. Three types of explanation

Resolving the global/local ambiguity requires distinguishing between (1) *global how-possibly*, (2) *local how-possibly*, and (3) *how-actually* explanations. Each type answers a distinct sort of question. The first type answers the question: could some potential process produce evolutionary changes in idealized populations? The second: could some known process produce, in a way consistent with the local information set for a real population, an observed evolutionary outcome or pattern? And the third: why, exactly, did some target evolutionary outcome or pattern occur?

To clarify the difference between these three types of explanation, I will represent them as different types of deductive explanatory argument along the lines of Railton (1978). To achieve generality, I have formulated them such that evolutionary processes produce outcomes with some non-zero probability; in the deterministic case the probability will be 1. Also, the production relation should be understood as a causal relation—natural selection, for example, produces change due to differences in fitness. For this reason I will adopt the following principle. If it is possible that a population has some set of properties, and there is an evolutionary model that identifies a possible process that will (causally) produce some outcome with a probability greater than zero given the property set, then that outcome is also possible. Call this the *causal principle of possibility*. Finally, evolutionary investigation need not follow the structure of these explanatory arguments. In fact, it will often deviate, as much of evolutionary inquiry seeks to reconstruct the original state of the population from the observed evolutionary outcomes. Nor must evolutionary explanations be arguments.⁴ I will consider each of the three types of explanation in turn.

Formal evolutionary models and theoretical investigation provide global how-possibly explanations. They have the form:

1a It is possible (relative to the global information set) that an idealized population p has some property set F .

Table 1

The relative fitness structure for the one-locus, two-allele system with additive fitness effects. The positive selection coefficient s determines the strength of selection against the A_2 allele. Let A_1 denote the fitter allele where the initial frequency of $A_1 = p$, and $A_2 = 1 - p = q$. Also, s is constrained such that no genotype fitness is less than zero; in this case let $0.5 \geq s > 0$.

Genotype	A_1A_1	A_1A_2	A_2A_2
Fitness	1	$1 - s$	$1 - 2s$
Frequency	p^2	$2pq$	q^2

1b For all populations x , if x has F then evolutionary process E will produce G in x with probability Y ($Y > 0$).

1c (The causal principle of possibility.)

\therefore It is possible (relative to the global information set) that the idealized population p exhibits outcome G produced by process E due to F .

These explanations focus primarily on formal or mathematical explication of premise 1b of the argument and relax biological restrictions on the information set assumed in premise 1a. This type of explanation neatly captures Lewontin's characterization of one investigative program of theoretical populations genetics: the 'minimal deductive program' that seeks 'to provide a rigorous network of relationships between the causal forces and their outcomes at the genetic level', and produces 'purely analytic results of an "if, then" form that can be used to demarcate the allowable from the unallowable claims of explanation' (Lewontin, 2000, p. 199). In simple Mendelian populations the models show that existing alleles can change into a new allele by the process of mutation. Random drift can cause allele frequencies to fluctuate in finite populations. Migration can introduce new alleles and change allele frequencies. Natural selection, due to differences in fitness between the alleles, also can cause evolutionary change. These processes, and only these processes, can change the genetic composition of the population. Notice that these idealized Mendelian systems are drastically simplified, ignoring many, often significant, details about molecular transmission, linkage, and demography. Although we loosely construct these idealized systems based on what we take to be general features of transmission genetics, there are no strong empirical restrictions on the formal models.

A simple example makes this clear. When there are fitness differences directional selection can change allele frequencies. Let us accept the standard idealizations that include random mating (Hardy-Weinberg ratios apply), additive fitness effects (no dominance), and an indefinitely large population (no significant drift occurs). The fitness structure for this model is given in Table 1. The model of directional selection specifies how the frequency of A_1 should change in a given generation (equation 1).

$$\Delta p = pqs / (1 - 2pqs - 2q^2s) \quad (1)$$

The outcome of this selection model is a consequence of mathematics. The frequency of A_1 will increase by Δp in the next generation, given the assumptions. The frequency will continue to change by Δp , a quantity determined by the parameter values of the previous

³ It is tempting to use Sober's (1984, p. 50) distinction between empirical *source laws* and mathematical *consequence laws* to describe the difference between local and global biological possibility in the following way. The local information set includes both the source and consequence laws, whereas the global information set includes just the consequence laws. This, however, will not work for two reasons. First, not all evolutionary models are mathematical. For example, laboratory models guide experimental design, as they have done with some group selection experiments (Griesemer & Wade, 1988). Darwin (1964 [1859]) originally formulated natural selection without mathematics, and there still are a number of useful non-mathematical descriptions of natural selection (Godfrey-Smith, 2009). Second, more than pure mathematics *alone* is needed to constrain biological possibility. The models must be models of evolutionary systems. Biologists build models, mathematical or otherwise, to shed light on biological phenomena. The motivation for constructing particular evolutionary models depends on what biologists take to be salient features of the biological world.

⁴ Salmon (1984, 1989) challenges the view that explanations are arguments on the grounds that arguments tend to be far less restrictive than explanations. They permit irrelevant premises and do not capture causal asymmetry, for example. I agree and so have included a causal production condition. Yet different arguments can help identify the differences between kinds of explanation even though they do not constitute explanations, hence my use of them here.

generation, until the allele fixes at a frequency of 1. The model of random drift, by suspending the assumption of an indefinitely large population and applying the mathematics of binomial sampling, provides another global how-possibly explanation. These models do not specify exactly how the actual course of evolution in target biological populations must unfold, although they do provide some constraints. Also, the models contain implicit universal quantifiers: they apply to all idealized Mendelian populations and hold for all permissible values of the parameters p , q , s .

Local how-possibly explanations aim to establish whether evolutionary processes can produce specific outcomes in some target biological system. They have the form:

2a It is possible (relative to the local information set) that a real target population r has some property set F .

2b For all populations x , if x has F then evolutionary process E will produce G in x with probability Y ($Y > 0$).

2c (The causal principle of possibility.)

\therefore It is possible (relative to the local information set) that the real population r exhibits outcome G produced by process E due to F .

These explanations focus primarily on empirical speculation about premise 2a. This speculation is constrained by the relevant local information set: the global information set determined by the theory plus the background biology of the target population r . Specifying premise 2b involves choosing a formal model with assumptions that approximate the target system and instantiating the parameters in the general evolutionary model with precise values.

Suppose Kettlewell sought first to explain how directional selection could possibly change the frequencies of light and dark morphs in the *Biston betularia* population. He would give a local how-possibly explanation that speculated about the biological parameter values (p , s) and whether the target system (the population of moths) approximated the model assumptions (random mating, additive fitness effects, large population size) necessary to make such an outcome (locally) possible. From what Kettlewell knew about moths and their environment, it is possible that melanism provides some additional camouflage from predation by birds, and that the trait approximates the assumed Mendelian genetic structure. There could be two alleles at one locus with the new allele for melanism (A_1) having additive effects. Perhaps having one copy of A_1 provides an intermediate degree of camouflage. These genes could have had consistent effects on fitness; moths with one or two copies of A_1 may have enjoyed a degree of reproductive success greater than those with two A_2 alleles due to greater camouflage from predators. Also, it is possible (relative to the local information set) that the population of moths approximates random mating and that the genes are transmitted from parents to offspring in a reliably Mendelian way. Additional local how-possibly explanations would identify other rival hypotheses for explaining the evolution of melanism.

Selecting the how-actually explanation from among the (local) biological possibilities requires evidence that supports the explanation over the other rival hypotheses. In contrast to how-possibly explanations, they have the form:

3a It is actually the case that the target population r has some property set F .

3b For all populations x , if x has F then evolutionary process E will produce G in x with probability Y ($Y > 0$).

3c (And process E actually produced G in r due to F)⁵

\therefore It is actually the case that the real population r exhibits outcome G produced by process E due to F .

The local how-possibly explanation uses the model of directional selection in a Mendelian population to make a prediction about how the genetic composition of the population of moths would change given the empirical speculations. This is one possible way for melanism to evolve. For Kettlewell to fashion a how-actually explanation he needs evidence about the predation selection pressures on light and dark morphs, the genetic architecture and demographics of moths, and the evolutionary outcome. By identifying confirmation-theoretic support relations between the set of local possibilities and the observations we can, in principle, determine which local possibility counts as the actual explanation for the target evolutionary outcome. The sort of observations that would count as evidence in the melanism case should discriminate between the selection-for-crypsis local possibility and the other local biological possibilities, such as (say) selection for another trait, perhaps metabolic efficiency, linked to melanism. Observations that dark moths do not outcompete light moths in environments with no predation would provide evidence for the selection-for-crypsis biological possibility over the linkage possibility (Kettlewell, 1955, pp. 324, 339).

The global how-possibly explanations have theory, mathematics, simulations, and analytical techniques as the resources for fashioning such explanations. They identify the biological possibilities relative to information about the capabilities and limitations of evolutionary processes for producing outcomes under different abstract assumptions (the global information set). The local how-possibly explanations draw upon the models of evolutionary processes and go one step further. They speculate about the biological possibilities relative to an information set enriched by the specific biology of a target system (the local information set). How-actually explanations, carefully confirmed by empirical tests, aim to identify the correct evolutionary processes that did, in fact, produce the target outcome.

4. Brandon on how-possibly versus how-actually

Brandon (1990, pp. 176–184) applies the distinction between how-possibly and how-actually explanations to evolutionary biology in order to distinguish incomplete from complete adaptation explanations. The difference between how-possibly and how-actually explanations is, on Brandon's view, one of degree rather than kind. And this seems to be the prevailing view. Resnik (1991, p. 143) follows Brandon, arguing that 'the difference between how-possibly and how-actually explanations is quantitative—a difference of degree—since empirical support comes in degrees'. Lewontin (2000, p. 198) identifies the minimal deductive program of population genetics as one extreme on a spectrum of investigative programs, and claims that such programs 'vary in the degree of their ambition, from the maximal inferential program meant to give a correct biological explanation of any and all observed evolutionary differentiation, down to a minimal deductive program that provides rules for recognizing acceptable and unacceptable explanations, without reference to any particular observed case'. Although Lewontin never uses such terms, this spectrum clearly

⁵ Railton (1978) makes the case that (how-actually) explanatory arguments about probabilistic phenomena need to include this 'parenthetical addendum'. I will follow Railton here. Evolutionary how-actually explanations explain why outcomes and patterns have actually occurred by uncovering the processes, chance or otherwise, that produced them. Premise 3b cites an evolutionary model to show that process E produces G_x from F_x with probability Y . Yet in the actual case either G_r occurs or it does not. Thus, Railton's parenthetical addendum provides 'information that is relevant to the causal origin of the explanandum' by communicating that it is the realization of some probabilistic process (ibid., p. 217). Premise 3c provides another way of distinguishing how-actually from how-possibly explanations in the probabilistic case.

describes a contrast between how-possibly and how-actually explanations.

While measuring dimensions of epistemological significance, such as the completeness of an adaptation explanation or the strength of empirical support, is certainly important, using the how-possibly/how-actually contrast as the metric distorts the role how-possibly explanations are supposed to play. Such a move trades one explanatory problem for another, to detrimental effect. The two separate problems are what Sober (2003) calls Hempel's problem and Peirce's problem. Hempel's problem concerns whether there are different *types* of explanation, and what counts as an *ideally complete* explanation type. Peirce's problem concerns the comparative *evidential support* for competing explanations, and which one counts as the *best* explanation given the evidence. Brandon and others focus on Peirce's problem: all adaptation explanations with less than ideal evidential support count as how-possibly explanations. Yet using the distinction to solve Peirce's problem relinquishes the resources to recognize differences between explanation types. Both global and local how-possibly explanations would merely mark intermediate points on the continuum of explanatory completeness with how-actually adaptation explanation at one extreme. In contrast, I take the how-possibly/how-actually distinction to apply better to Hempel's problem: delineating different types of explanation.

For Brandon a how-actually explanation includes a suite of evidence that makes it an 'ideally complete adaptation explanation' (Brandon, 1990, p. 165). These complete explanations include evidence: (1) that selection has occurred; (2) for the ecological basis of selection; (3) that the trait is heritable; (4) of population structure; and, (5) of the ancestral character states. Brandon argues that all five kinds of evidence are necessary to establish a how-actually explanation. Notice that this entails that there are very few how-actually explanations to be found in evolutionary biology, and that all such explanations that exist are necessarily well confirmed. This is an odd consequence since most of evolutionary inquiry fails to produce ideally complete explanations, and even the best explanations available will have varying evidential support for each component. Brandon's how-possibly explanations, on the other hand, lack one or more of the components of an ideally complete adaptation explanation. These explanations have an intermediate degree of empirical support (ibid., p. 184), but are still useful because they provide part of the complete explanation of how selection caused particular evolutionary outcomes.

While Brandon provides a good metric of evidential support for adaptation explanations, partial explanations are not properly described as how-possibly explanations. Recall that Dray (1957, pp. 156–159) formulates how-possibly explanation as aiming at a different explanatory goal than the (how-actually) explanation type that concerned Hempel (1965): they aim to show that some outcome is possible rather than why it actually occurred. Brandon (1990, p. 178 n. 18) explicitly formulates how-possibly explanations differently, using an example about whether tropical plants could possibly grow in Switzerland as motivation. A how-possibly explanation should, according to Brandon, show how these plants actually got to Switzerland as well as how it would be possible to cultivate plants in Switzerland's climate. Distinguishing the global and local questions about biological possibility resolves the example in finer detail. A how-possibly explanation of how tropical plants may manage to survive in temperate climates and how they may migrate to such climes provides an answer to the global question (for example, climate models show that mountain range shadows make local climates milder and organisms can have huge dispersal ranges), whereas a how-possibly explanation of how these particular tropical plants could appear on the slopes of the Swiss Alps answers the local question (for example, horticulturists may have brought tropical plants to Switzerland and cultivated

them in a local range shadow). With *evidence* that horticulturists indeed imported and cultivated tropical plants, and that the Swiss Alps create a range shadow that provides a milder climate, the local how-possibly explanation becomes a how-actually explanation, though perhaps incompletely confirmed. What Brandon counts as how-possibly explanations should, instead, count as incompletely confirmed or tentative how-actually explanations because they provide partial explanations with some empirical support for how evolution probably unfolded to produce an actual outcome.

Brandon's example of a purported how-possibly evolutionary explanation is the tentative explanation for how insect wings may have evolved by initial selection for thermoregulation followed by later modification by selection for flight (Kingsolver & Koehl, 1985). Kingsolver and Koehl, however, do not merely provide a local how-possibly explanation for the evolution of flight in insects, nor should they describe their project as such. They conduct manipulation experiments in the laboratory to determine the effects of wing morphologies for both aerodynamics and thermoregulation. Based on their results they conclude that they have evidence that suggests insect wings evolved first for thermoregulation and later were co-opted and evolved to function as wings for flight. The exact point where thermoregulatory wings were co-opted for flight depends on changes in body size in insect lineages (ibid., p. 503). Their study does not merely speculate about the possible evolutionary trajectories that could have produced insect wings, but provides evidence that selection for thermoregulation was followed by selection for flight. The tentative adaptation explanation for insect wings simply lacks perfect empirical support; it is missing one or more of the components of an ideally complete adaptation explanation.

Brandon wants his how-possibly explanations to have some evidential support because he seeks to distinguish them from the 'just-so' stories that Gould and Lewontin (1979) criticize. He claims that, unlike just-so stories, his how-possibly explanations 'can be rigorously formulated so that they do have testable consequences' (Brandon, 1990, p. 183), and have some intermediate degree of completeness or evidential support. But changing the problem, from Hempel's problem of identifying different types of explanation to Peirce's problem of assessing comparative evidential support, does not help clarify the role of just-so evolutionary stories. In the limit, with no supporting evidence, a potential adaptation explanation would count, on my view, as a local how-possibly explanation. And local how-possibly explanations *are* just-so stories that speculate about the adaptive (or non-adaptive) evolutionary history of a lineage. There is nothing wrong with telling just-so stories so long as they are taken for what they are: *speculations* about whether some evolutionary process makes an outcome or pattern possible (Sterelny & Griffiths, 1999, p. 225). The practice Gould and Lewontin deplore is the unceasing *ad hoc* revision of adaptation just-so stories in light of recalcitrant evidence, driven by the assumption that most traits are adaptations (Gould & Lewontin, 1979, p. 586).

Later Brandon describes how-possibly explanations differently and in a way consistent with the global type: 'I view much of the theoretical work in mathematical population genetics as the construction and testing of how-possibly explanations (consider kin and group selection models for the evolution of altruism)' (Brandon, 1990, p. 184). This sort of theoretical exploration does provide how-possibly explanations of the global type, for it relies on formal inquiry and addresses an abstract problem (the evolution of altruism) detached from any specific biological system. This is different than the incomplete evolutionary explanations Brandon uses as examples of explaining how possible. Treating purely theoretical work on group selection and incomplete empirical adaptation explanations as the same kind of explanation confounds the global and local questions about biological possibility. The grouping also

equates explanatory success with empirical support. These differences should be respected, if we hope to capture the rich explanatory practices of evolutionary biology.

Yet Brandon's claim that *both* count as the same type of explanation, and that both are relevant to Peirce's problem, raises an intriguing question about the nature of confirmation. Can formal, theoretical inquiry test evolutionary models and provide confirmation-theoretic support for them? This question concerns the scope and power of how-possibly explanation, for it asks whether such explanations can provide evidence for their explananda. Although there is a relationship between how-possibly explanation and confirmation, this is not it, and it will clarify my account to explore why.

5. Formal sources of confirmation?

Abstract, usually mathematical models play a crucial and multifaceted role in evolutionary theory. This has led some to suggest that the formal analysis of these complex models can provide a sort of confirmation that the models truly or adequately represent the biological world. Levins (1966, 1993), and later Wimsatt (1981), defend *robustness analysis*—a procedure where modelers determine whether various idealized models with subtly different assumptions predict the same outcomes or processes—as a way of providing evidence for the resulting 'robust theorems'. Orzack and Sober (1993) object that robustness analysis, a sort of mathematical or logical inquiry, cannot provide any kind of confirmation; it can only reveal novel deductive consequences of models we already take to be true. Weisberg (2006) replies that robustness analysis helps determine whether idealized models of complex systems identify and represent robust causal connections, and so provides a kind of 'low-level confirmation'. Plutynski (2001, p. S229) makes a similar argument about simulations in evolutionary theory: 'mathematical or computer simulation models can function as tests of theoretical questions in the same way as do laboratory experiments'. Dietrich (1996) takes the opposite view and argues that simulations do not count as true tests because they do not compare the models to any biological data. I will not attempt to resolve all the rich philosophical issues about the nature of idealization, representation, and evolutionary theorizing that emerge in these debates. Instead, I will focus on the nature of confirmation, and use the distinction between global and local questions to provide a novel perspective on what this formal inquiry does and why it appears to have confirmatory power.

The formal inquiry exemplified by robustness analysis and simulation provides global how-possibly explanations that constrain what counts as a biological possibility. Recall that global how-possibly explanations identify the potential evolutionary processes and show how these abstract processes produce all sorts of outcomes. Simulations and robustness analysis do just this. Simulations explore the capabilities of complex models by repeatedly running a process model on a computer using different values for the different parameters and observing the kinds of outcomes that occur. Only the constraints introduced by very general assumptions about the evolutionary process (that is, the global information set) matter. The formal investigation is unconstrained by biological information about any target population.

To briefly illustrate how formal inquiry constrains possibility, consider a simple example of an evolutionary simulation. Using the model of directional selection with the fitness structure specified by Table 1 (Sect. 2), one can simulate how selection acts in idealized populations. A simulation involves specifying various sets of values for the parameters (p , s) and iterating equation 1 to determine how p changes each generation over time. No information about any target population constrains what values the parameters

take. Simulations use many different parameter values to reveal different possible evolutionary trajectories for the idealized Mendelian population. By determining what is possible, the simulation also determines what trajectories are impossible for the directional selection model to produce. They show that for constant, non-zero values of s there can be no equilibrium where both A_1 and A_2 alleles are maintained in the population. Mathematical exploration of the evolutionary model through simulation carves up possibility space and so constrains what counts as biologically possible (relative to the global information set).

On the issue of robustness analysis, Weisberg (2006, p. 738) gives a clear reconstruction of the analytical technique as a procedure that identifies robust theorems of the form: '*ceteris paribus*, if [common causal structure] obtains, then [robust property] will obtain', where [common causal structure] and [robust property] are parts of the robust theorem. Weisberg offers an excellent example of robustness analysis based on the Lotka–Volterra predator–prey models. The formal analysis supposedly provides 'low-level confirmation' for the robust theorem. Weisberg agrees that this 'low-level' kind of confirmation differs from standard direct confirmation of theory by evidence, but nonetheless counts as a kind of confirmation-theoretic support. Specifically, low-level confirmation provides 'confirmation of the fact that certain mathematical structures can adequately represent properties of target phenomena' (ibid., p. 740).

This sort of formal inquiry exemplifies a global how-possibly explanation. Why should we also take it as providing evidence for or tests of the evolutionary models or robust theorems? Robustness analysis explores the connections between model structures—freely varying initial conditions and parameters that Weisberg takes to represent the causal structure—and kinds of evolutionary outcomes described by the robust property. This sort of inquiry provides 'if-then' connections much like those that Lewontin identifies as part of the minimal deductive program of population genetics. Insofar as the formal inquiry into models establishes possible connections between model structures and the robust property, it constrains what counts as biologically possible. *Evidence* for a robust theorem should do more than constrain possibility space; it should help show that the model structures are instantiated in *actual biological systems*. I agree with Weisberg that something sets robust models apart. Perhaps robust models tend to better at representing biological systems, or perhaps they provide a better mathematical framework for identifying causal rather than spurious regularities. But such a difference between robust and fragile models does not turn on robustness analysis conferring extra confirmation-theoretic support. Robustness analysis can provide support for the *representational adequacy* of a model in the way Weisberg envisions, but this is simply not the same as providing confirmation that the model *actually represents* causal relationships in a target system.

Formal inquiry into evolutionary models, a kind of indirect or non-empirical inquiry, constrains the space of biological possibility. Confirmation, the product of direct *empirical* inquiry, narrows down which possibility in the space is actually the case by supporting certain biological possibilities over others. Formal inquiry indeed plays a central role in evolutionary science, but equating it with testing and confirmation obscures this role. Robustness analysis does not provide confirmation-theoretic support to robust theorems, and simulations do not provide tests in the same way as empirical experiments. With respect to simulations, Dietrich provides a better analysis of their role: simulations assess the 'computational adequacy' of models rather than their empirical adequacy since they compare different types of models rather than the model and the world (Dietrich, 1996, pp. 347–350). While both formal inquiry (robustness analysis, simulation) into idealized systems and empirical inquiry into real systems (testing, confirmation)

aim to illuminate the evolutionary processes in the biological world, they differ in the strategies deployed to realize this aim.

The difference is subtle but important. It is subtle because both aim to narrow down possibility space. Formal inquiry helps determine the biological possibilities (relative to the global information set), and confirmation helps identify which possible rival (relative to a specific local information set) actually explains some target evolutionary outcome. This similarity makes it seem that both sorts of inquiry provide some ‘evidence’ that certain biological possibilities are actually the case, and therefore that formal inquiry yields a kind of confirmation-theoretic support. The difference is important because the strategies for narrowing down possibility space incorporate biological information differently. Formal inquiry explores the global scope of biological possibility by abstracting or idealizing away from the details of any target system, whereas empirical inquiry explores the local scope of possibility within the boundaries of specific biological systems, and with enough supporting evidence determines what is actual. Only the product of latter kind of inquiry, the collected data that discriminate between locally possible explanations for target phenomena, can provide *confirmation* of evolutionary hypotheses. Confirmation deals with the theory–data relation, and so we should restrict its application to direct empirical inquiry into the target systems that make up the biological world. While formal inquiry plays a crucial role in evolutionary biology by structuring and constraining the space of biological possibility, a role that deserves due philosophical attention, we should resist describing this kind of inquiry as confirmation-theoretic, or risk obscuring what is distinctive about each form of inquiry.

6. Expanding the scope of biological possibility

Any theory of confirmation needs a principled account of what possibilities deserve consideration, for evidential relations depend on facts about all hypotheses under consideration.⁶ How-possibly explanation provides such an account by identifying the genuine rivals: what we take to be the global and local biological possibilities. To support my argument that how-possibly explanation is relevant to confirmation, I will marshal an example from the rise of the neutral theory of molecular evolution (NTME).⁷ I will focus on one important argument, one that helped NTME facilitate a transition from a state where hypotheses of random drift had a minor explanatory role to a new state where drift became a genuine rival for explaining evolutionary change at the molecular level. To argue that the NTME example engendered such a shift in what we take to be possible, I will review the consensus within a prevailing theoretical perspective, and identify a global how-possibly explanation that helped establish drift as a genuine rival.⁸ Examining the transition shows that new rivals, introduced by successful how-possibly explanation, alter subsequent testing practice.

The neutral theory emerged in the context of views about the relevant evolutionary alternatives held by prominent evolutionists, such as Mayr, Dobzhansky and Simpson. Simpson (1964) provides a clear statement of the consensus. In an article responding to the increasing interest in biological macromolecules (protein and DNA), Simpson argues that DNA cannot be completely responsible for an organism’s phenotype (*ibid.*, p. 1536), that there exist few, if any, alleles without adaptive or deleterious consequences of some kind (*ibid.*, p. 1537), and that molecular data is unreliable, espe-

cially for determining phylogenetic relationships and estimating the time since divergence (a quantity important for the molecular clock) (*ibid.*, pp. 1535–1537). Thus, Simpson argued that the pan-selectionism that characterized morphological evolution should extend to the molecular level (Dietrich, 1994, pp. 21–22).

The only class of drift hypotheses they considered a serious possibility involved founder effects. Founder effects occur when a relatively small subpopulation, perhaps a few immigrants arriving on an island, found a new population of organisms. Because the founding population is small only a random subset of the total genetic variation in the original population gets included in the new population. The random sampling of genetic variation counts as a type of drift, and different subsets of variation may have different consequences for adaptive evolution in that new population. Dobzhansky and Pavlovsky (1957, p. 315) found that small founding populations can significantly increase the variability in the outcomes of subsequent selection. Mayr’s (1954, 1963) work on island biogeography provides evidence that such founder effects occur in natural populations and can lead to speciation. Otherwise Mayr and others remained staunchly committed to selection for explaining evolution (see, for example, Mayr, 1983). In their view drift simply lacked the power to effect significant evolutionary change.

The proposal of NTME by Kimura (1968) and King and Jukes (1969), however, challenged the consensus by rehabilitating a new class of drift hypotheses. The key move of NTME was to treat most, if not all, detectable molecular changes as *neutral*, and thus having no effect on fitness. These neutral mutations change frequency in populations through genetic drift rather than natural selection. On their proposal selection still acts, but in a purifying way by removing deleterious mutations from the population. Any observed variation that slipped through the sieve of selection would be neutral. Notice that NTME did not introduce any new evolutionary processes. It used the old theoretical resources in a novel way, and (among other things) constructed a global how-possibly explanation for a general evolutionary pattern: a constant and fast rate of molecular evolution.

The general pattern became a salient explanandum due to the ‘molecularization’ of biology (Dietrich, 1994; Suarez & Barahona, 1996). New experimental techniques introduced an immense amount of molecular data. Biologists started using gel electrophoresis to detect variation in proteins (Hubby & Lewontin, 1966; Lewontin & Hubby, 1966), and additional techniques in biochemistry provided more molecular data on the sequences and functionality of proteins (King & Jukes, 1969). As part of this transition, Zuckerkandl and Pauling (1965) proposed the *molecular clock* hypothesis based on data that showed a linear relationship between the number of changes accumulating in various proteins and divergence time. This comparative data suggested that the rate of protein evolution is *constant* over evolutionary time, hence the proposal that there exists a kind of molecular clock. Given that selection depends on relatively ephemeral ecological factors of the environment, the rate of protein evolution should be erratic rather than constant, if selection played the important role.

NTME explained how a constant rate of molecular evolution produced solely by random drift is possible. The explanatory argument goes like this. Suppose most mutations are neutral and therefore that genetic drift is the primary evolutionary process. Suppose also that the mutation rate remains relatively constant over evolutionary time. These assumptions are consistent with the general

⁶ Formal Bayesianism makes this clear. Whether, and the degree to which, an observation confirms or disconfirms a particular hypothesis depends on the likelihoods of all hypotheses under consideration (see n. 9).

⁷ The origin of NTME has a rich historical explanation. For detailed exploration see Lewontin (1974), Dietrich (1994, 1998), Ohta & Gillespie (1996), Suarez & Barahona (1996).

⁸ While I will investigate the perspective that created much of the resistance to the initial proposal of NTME, there were different perspectives that took drift more seriously. For example, Wright accorded drift an important role in his shifting balance theory of evolution (Provine, 1986; Skipper, 2002), and Lamotte argued that drift explained phenotypic evolution in *Cepaea nemoralis* (Millstein, 2007).

assumptions of population genetics (that is, they do not contradict the global information set). Kimura (1968) showed that drift alone produced a rate of evolution (k) equal to the rate of mutation (μ):

$$k = \mu \quad (2)$$

Kimura and Ohta (1971) provide the clearest analysis of equation 2. They show that the rate of fixation equals the product of three quantities: the probability of a mutation occurring in a particular allele (the rate of mutation μ), the number of alleles in a diploid population (twice the population size N), and the probability that one of these new mutations will fix. According to Kimura's (1962) model, the probability of fixation for a neutral allele in generation t is equal to the frequency of the allele at t , so the probability of fixation for a new mutation equals one over the total number of alleles ($1/2N$).

$$k = (\mu) \cdot (2N) \cdot (1/2N) \quad (3)$$

The quantity $2N$ cancels and equation 3 reduces to equation 2. Notice that the rate of fixation is independent of population size. Random drift can thus produce a constant and significant rate of evolution in any population, regardless of size. NTME generated a global how-possibly explanation for the clock-like trend in the molecular data. The explanation above counts as a how-possibly one because it does not marshal any evidence for the actual rate of evolution. The how-actually explanations for real molecular evolutionary rates remain controversial (Ohta & Gillespie, 1996; Cutler, 2000). It also dispelled the presumption of the prevailing theoretical perspective that drift could not produce significant evolutionary change. The explanation counts as a *global* how-possibly one because it appeals to theoretical models to show how a constant rate of molecular evolution is consistent with background formal assumptions. Such models must be applied to a target system to explain specific rates of molecular evolution in real lineages.

The global how-possibly explanation changed the way biologists test hypotheses about molecular evolution. Given that drift can produce significant molecular evolution, tests for selection at the molecular level must consider the neutral rival. The most successful tests for adaptive molecular evolution, such as the HKA and MK tests, contrast neutral models with positive (directional) selection models. In an authoritative review of methods for detecting selection at the molecular level, Kreitman (2000, pp. 541–542) claims that the NTME 'is the backbone for evolutionary analysis of DNA sequence variation and change for three reasons': (1) much of the genome is non-functional, and so the neutral model captures the evolution of these regions; (2) selection at linked sites can affect the evolution of neutral mutations, and so deviations from the neutral model can be informative; and (3) 'selective neutrality is a useful null hypothesis against which to test for evidence of selection'. Crow (1987) also argues that the neutral model is the proper null model for evolutionary inquiry.

Confirmation relations depend on the nature of rival hypotheses considered. This claim almost has the status of a truism in confirmation theory. In Bayesianism, for example, whether some evidence increases our credence in a hypothesis (that is, when E confirms H) depends on the likelihoods of all hypotheses under consideration.⁹ A how-possibly explanation, in effect, identifies a new rival and helps determine the likelihood for that rival. Thus, explaining how possible can have a vast impact on confirmation by requiring alterations and innovations in testing methods. The global how-possibly explanation provided by NTME illustrates how profound and immediate this effect can be in scientific practice.

7. Conclusion

Any theory of confirmation must have an account of what marks out the rival hypotheses from the rest, for not all conceivable hypotheses merit consideration. Explaining how possible informs confirmation by constraining (what we take to be) the biological possibilities. To shed philosophical light on explaining how possible, I have given an account that distinguishes three types of explanation: global how-possibly, local how-possibly, and how-actually.

The contrast between how-possibly and how-actually explanation is familiar, but often misdescribed as a difference in the degree of completeness or empirical support. The contrast is best understood as a difference in type. To successfully explain the actual course of evolution we must have sufficient evidence. Yet we can successfully explain how possible without *any* empirical support; we need only show that some outcome or pattern is consistent with a specific set of information.

The distinction between global and local constraints on biological possibility is new. This attempts to capture the difference between two sorts of inquiry into what is biologically possible: general, theoretical exploration of abstract models versus targeted empirical speculation, testing, and application of those models to real biological systems. Many try to unify these disparate types of inquiry, and even argue that formal inquiry provides a kind of confirmation-theoretic support for evolutionary models. This is a mistake. Confirmation-theoretic support comes from connecting models to the world, rather than assessing the models themselves. Treating formal inquiry into what is possible (relative to the global information set) and empirical inquiry into what is possible (relative to the local information set for the target biological system) as the same type of how-possibly explanation obscures the differences between investigating models versus investigating the world.

The origin of NTME provides an example, one of many, from scientific practice where explaining how possible has important consequences for confirmation. Articulating how neutral evolution can produce a fast and constant rate of molecular evolution altered testing practice. The neutral model is now an essential rival for any molecular test for selection. Confirmation formalizes how evidence supports our conjectures. And assessing the evidence depends on what we take to be possible.

Acknowledgements

This paper traces its origin back to my dissertation work, and so I owe many thanks to the community of philosophers at Stanford and ANU for the support and input necessary to make these ideas work. And for detailed comments, often on several drafts, I owe thanks to Elliott Sober, Peter Godfrey-Smith, Kim Sterelny, Michael Strevens, Stephen Downes, Samir Okasha, Michael Weisberg, Roberta Millstein, Robert Skipper, Derek Turner, and the anonymous referees.

References

- Beatty, J. (1995). The evolutionary contingency thesis. In G. Wolters, & J. G. Lennox (Eds.), *Concepts, theories, and rationality in the biological sciences* (pp. 45–81). Pittsburgh: University of Pittsburgh Press.
- Beatty, J. (1997). Why do biologists argue like they do? *Philosophy of Science*, 64(Suppl.), S432–S443.
- Brandon, R. (1990). *Adaptation and environment*. Princeton, NJ: Princeton University Press.

⁹ Here are the technical details. E confirms H iff $P(H|E) > P(H)$ where $P(H|E) = (P(E|H)P(H))/P(E)$. Note that $P(E) = P(E|H)P(H) + P(E|\neg H)P(\neg H)$, and the quantity $P(E|\neg H)$ depends on the likelihoods of the other hypotheses under consideration. See, for example, Earman (1992), Howson & Urbach (1993).

- Crow, J. F. (1987). Neutral models in molecular evolution. In M. H. Nitecki & A. Hoffman (Eds.), *Neutral models in biology* (pp. 11–22). Oxford: Oxford University Press.
- Cutler, D. J. (2000). Understanding the overdispersed molecular clock. *Genetics*, 154(3), 1403–1417.
- Darwin, C. (1964). *The origin of species* (facsim. ed.). Cambridge, MA: Harvard University Press. (First published 1859)
- Dawkins, R. (1986). *The blind watchmaker*. New York: W. W. Norton and Company.
- Dietrich, M. (1994). The origins of the neutral theory of molecular evolution. *Journal of the History of Biology*, 27, 21–59.
- Dietrich, M. (1996). Monte Carlo experiments and the defense of diffusion models in molecular population genetics. *Biology and Philosophy*, 11, 339–356.
- Dietrich, M. (1998). Paradox and persuasion: Negotiating the place of molecular evolution within evolutionary biology. *Journal of the History of Biology*, 31, 85–111.
- Dobzhansky, T., & Pavlovsky, O. (1957). An experimental study of the interaction between genetic drift and natural selection. *Evolution*, 11(3), 311–319.
- Dray, W. (1957). *Law and explanation in history*. Oxford: Oxford University Press.
- Dray, W. (1993). *Philosophy of history* (2nd ed.). Englewood, NJ: Prentice Hall.
- Earman, J. (1992). *Bayes or bust? A critical examination of Bayesian confirmation theory*. Cambridge, MA: MIT Press.
- Godfrey-Smith, P. (2009). *Darwinian populations and natural selection*. New York: Oxford University Press.
- Gould, S. J. (1989). *Wonderful life: The Burgess Shale and the nature of history*. New York: W. W. Norton.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society London, Series B, Biological Sciences*, 205, 581–598.
- Griesemer, J. R., & Wade, M. J. (1988). Laboratory models, causal explanations and group selection. *Biology and Philosophy*, 3, 67–96.
- Hagen, J. (1999). Retelling experiments: H. B. D. Kettlewell's studies of industrial melanism in peppered moths. *Biology and Philosophy*, 14, 39–54.
- Hempel, C. G. (1965). *Aspects of scientific explanation and other essays in the philosophy of science*. New York: Free Press.
- Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach* (2nd ed.). Chicago: Open Court.
- Hubby, J. L., & Lewontin, R. C. (1966). A molecular approach to the study of genic heterozygosity in natural populations. I. the number of alleles at different loci in *Drosophila pseudoobscura*. *Genetics*, 54, 577–594.
- Jackson, F., & Pettit, P. (1992). In defense of explanatory ecumenism. *Economics and Philosophy*, 8, 1–21.
- Kettlewell, H. (1955). Selection experiments on industrial melanism in the Lepidoptera. *Heredity*, 9, 323–342.
- Kettlewell, H. (1956). Further selection experiments on industrial melanism in the Lepidoptera. *Heredity*, 10, 287–301.
- Kimura, M. (1962). On the probability of fixation of mutant genes in populations. *Genetics*, 47, 713–719.
- Kimura, M. (1968). Evolutionary rate at the molecular level. *Nature*, 217, 624–626.
- Kimura, M., & Ohta, T. (1971). Protein polymorphisms as a phase of molecular evolution. *Nature*, 229, 467–469.
- King, J. L., & Jukes, T. H. (1969). Non-Darwinian evolution. *Science*, 164, 788–798.
- Kingsolver, J. G., & Koehl, M. A. R. (1985). Aerodynamics, thermoregulation, and the evolution of insect wings: Differential scaling and evolutionary change. *Evolution*, 39, 488–504.
- Kreitman, M. (2000). Methods to detect selection in populations with applications to humans. *Annual Review of Genomics and Human Genetics*, 1, 539–559.
- Levi, I. (1988). Four themes in statistical explanation. In B. Skyrms, & W. L. Harper (Eds.), *Causation in decision, belief change, and statistics*, Vol. 2 (pp. 195–222). Dordrecht: Kluwer Academic.
- Levins, R. (1966). The strategy of model building in population biology. *American Scientist*, 54, 421–431.
- Levins, R. (1993). A response to Orzack and Sober: Formal analysis and the fluidity of science. *Quarterly Review of Biology*, 68(4), 547–555.
- Lewontin, R. C. (1974). *The genetic basis of evolutionary change*. New York: Columbia University Press.
- Lewontin, R. C. (1985). Population genetics. In P. Greenwood, P. Harvery, & M. Slatkin (Eds.), *Evolution: Essays in honour of John Maynard Smith* (pp. 3–18). Cambridge: Cambridge University Press.
- Lewontin, R. C. (2000). What do population geneticists know and how do they know it? In R. Creath & J. Maienschein (Eds.), *Biology and epistemology* (pp. 191–214). Cambridge: Cambridge University Press.
- Lewontin, R. C., & Hubby, J. L. (1966). A molecular approach to the study of genic heterozygosity in natural populations. II. Amount of variation and degree of heterozygosity in natural populations of *Drosophila pseudoobscura*. *Genetics*, 54, 595–609.
- Mayr, E. (1954). Change of genetic environment and evolution. In J. Huxley, A. C. Hardy, & E. B. Ford (Eds.), *Evolution as a process* (pp. 157–180). London: Allen and Unwin.
- Mayr, E. (1963). *Animal species and evolution*. Cambridge, MA: Harvard University Press.
- Mayr, E. (1983). How to carry out the adaptationist program? *American Naturalist*, 121(3), 324–334.
- Millstein, R. (2007). Distinguishing drift and selection empirically: 'The great snail debate' of the 1950s. *Journal of the History of Biology*, 41, 339–367.
- Nilsson, D.-E., & Pelger, S. (1994). A pessimistic estimate of the time required for an eye to evolve. *Proceedings of the Royal Society London, Series B, Biological Sciences*, 256, 53–58.
- O'Hara, R. J. (1988). Homage to Clio, or, towards an historical philosophy for evolutionary biology. *Systematic Zoology*, 37(2), 142–155.
- Ohta, T., & Gillespie, J. H. (1996). Development of neutral and nearly neutral theories. *Theoretical Population Biology*, 49, 128–142.
- Orzack, S. H., & Sober, E. (1993). A critical assessment of Levins's the strategy of model building in population biology. *Quarterly Review of Biology*, 68(4), 533–546.
- Plutynski, A. (2001). Modeling evolution in theory and practice. *Philosophy of Science*, 68(Suppl.), S225–S236.
- Plutynski, A. (2004). Explanation in classical population genetics. *Philosophy of Science*, 71, 1201–1204.
- Plutynski, A. (2005). Explanatory unification and the early synthesis. *British Journal for the Philosophy of Science*, 56, 595–609.
- Provine, W. B. (1986). *Sewall Wright and evolutionary biology*. Chicago: University of Chicago Press.
- Railton, P. (1978). A deductive-nomological model of probabilistic explanation. *Philosophy of Science*, 45, 206–226.
- Reiner, R. (1993). Necessary conditions and explaining how-possibly. *The Philosophical Quarterly*, 43(170), 58–69.
- Resnik, D. B. (1991). How-possibly explanations in biology. *Acta Biotheoretica*, 39, 141–149.
- Rudge, D. (1999). Taking the peppered moths with a grain of salt. *Biology and Philosophy*, 14, 9–37.
- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton, NJ: Princeton University Press.
- Salmon, W. (1989). *Four decades of scientific explanation*. Minneapolis: University of Minnesota Press.
- Salmon, W. C. (1990). Rationality and objectivity in science or Tom Kuhn meets Tom Bayes. In W. Savage (Ed.), *Scientific theories*, Vol. 14 (pp. 175–204). Minneapolis: University of Minnesota Press.
- Simpson, G. G. (1964). Organisms and molecules in evolution. *Science*, 146(3651), 1535–1538.
- Skipper, R. A. (2002). The persistence of the R. A. Fisher–Sewall Wright controversy. *Biology and Philosophy*, 17, 341–367.
- Sober, E. (1984). *The nature of selection*. Chicago: University of Chicago Press.
- Sober, E. (1990). Contrastive empiricism. In W. Savage (Ed.), *Scientific theories*, Vol. 14 (pp. 392–412). Minneapolis: University of Minnesota Press.
- Sober, E. (2003). Two uses of unification. In F. Stadler (Ed.), *The Vienna Circle and logical empiricism: Re-evaluation and future prospects* (pp. 205–216). Dordrecht: Kluwer Academic Publishers.
- Sterelny, K. (1996). Explanatory pluralism in evolutionary biology. *Biology and Philosophy*, 11, 193–214.
- Sterelny, K., & Griffiths, P. E. (1999). *Sex and death: An introduction to philosophy of biology*. Chicago: University of Chicago Press.
- Suarez, E., & Barahona, A. (1996). The experimental roots of the neutral theory of molecular evolution. *History and Philosophy of the Life Sciences*, 18, 55–81.
- Van Fraassen, B. C. (1977). The only necessity is verbal necessity. *Journal of Philosophy*, 74, 71–85.
- Weisberg, M. (2006). Robustness analysis. *Philosophy of Science*, 73, 730–742.
- Wimsatt, W. C. (1981). Robustness, reliability, and overdetermination. In M. Brewer & B. Collins (Eds.), *Scientific inquiry and the social sciences* (pp. 124–163). San Francisco: Jossey-Boss.
- Wimsatt, W. C. (2002). Using false models to elaborate constraints on processes: Blending inheritance in organic and cultural evolution. *Philosophy of Science*, 69(Suppl.), S12–S24.
- Woodward, J. (2009). Scientific explanation. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/entries/scientific-explanation/>.
- Zuckerkindl, E., & Pauling, L. (1965). Evolutionary divergence and convergence in proteins. In V. Bryson & H. Vogel (Eds.), *Evolving genes and proteins* (pp. 97–166). New York: Academic Press.