

Teleological Explanation: Surveying the Landscape

Jonathan Birch

This paper presents a novel account of teleological explanations in biology. I outline the “shorthand approach” to such explanations, on which they are taken to convey implicit evolutionary explanations. “Selected effects” accounts of teleological explanation dominate recent literature, but they struggle to accommodate teleological explanations of complex traits built through cumulative selection. I articulate the general notion of a *landscape explanation*, which, applied to biology, explains the evolution of complex features in a population by citing salient features of the population’s fitness landscape. I show that such explanations lend themselves to a teleological shorthand. I close with remarks concerning when a teleological explanation of a trait is legitimate, and why teleological language strikes us as appropriate when we give evolutionary explanations.

1. The Shorthand Approach

Why do sparrows have wings? *For flying*. Why does the heart beat? *In order to circulate the blood*. Teleological explanations often strike us as apt in biological contexts. Such explanations purport to explain the presence of some trait in a type of organism by citing an *effect* the trait produces. Informally, they explain why a trait is *there* by stating what it is *there for*. In so doing, they often make use of language otherwise reserved for artefacts and agents: phrases like “in order to”, and terms like terms “purpose”, “goal” and “function”.

Philosophical responses to such explanations fall into two broad camps. One aims to bring teleological explanations within the scope of broader theories of scientific explanation (e.g. Ayala 1970, Wright 1976, Nagel 1977, Salmon 1989, Brandon 1990, Neander 1991a,b). The other seeks to deny such explanations scientific legitimacy (e.g. Hempel 1959, van der Steen 1971, Ghiselin 1974, 1994, Bechtel 1989, Cummins 1975, 2002, Kramer 1984, Hanke 2004). The motive for adopting the latter position is clear. We often give teleological explanations that pertain to the features of *artefacts*: my computer has a mouse for moving the cursor, my kettle has a heating element for boiling water. But these explanations work on the implicit presupposition that these objects have been *designed* with a particular purpose in mind. We can account for such explanations by realizing that they implicitly point to a cause of the feature they purport to explain: the cause is the designer's *intention* that the heating element boil water or that the mouse move the cursor. The worry is that teleological explanation in biology presupposes an implicit, radical commitment to the view that organisms too are products of intentional design. On the assumption that there is no strange backward causation at work, it is hard to see how else a trait's effects can explain why it is there.

How can the defender of the scientific legitimacy of teleological explanations dispel such doubts? One promising line of thought suggests that teleological explanations, in the context of modern biology, serve as a *shorthand* for evolutionary explanations. I will call this the "shorthand approach". The appeal of the shorthand approach lies in the hope that, if we can account for how the "longhand" evolutionary explanation allows a trait's effects to play a role in explaining its current presence, we will be able to account for how our everyday teleological explanations of heartbeats and kidneys can be scientifically legitimate after all.

But the shorthand approach needs fleshing out. Important questions remain: What is the evolutionary longhand implied by the teleological shorthand? And how does the longhand legitimize the explanation of a trait's current presence by appeal to its effects?

Karen Neander (1991a,b), influenced by Larry Wright (1973, 1976), supplies a theory of teleology in biology that purports to answer these questions. For Neander, teleological explanations cite "proper functions", where a trait's proper function is defined as that effect it produces in virtue of which it was selected. Robert Cummins (2002) raises a problem for this "selected effects" account of teleological explanation. In the next section I argue that, while Cummins's objection is not decisive, it presents a challenge to the "selected effects" theory serious enough to motivate a search for a different way to flesh out the shorthand approach. In the rest of the paper, I proceed to argue for such an alternative. I propose that teleological explanations work by drawing our attention to salient features of a population's *fitness landscape*.

2. The "Selected Effects" Theory

2.1 Wright and Neander

The "selected effects" (SE) theory derives from Wright's (1973) observation that, in populations under natural selection, a trait can come to have what Wright calls a *consequence etiology*, whereby its current presence is explained by a type of effect its past tokens produced. Suppose some trait *T* has been selected in virtue of producing some advantageous effect *E*; and suppose that, had *T* not been selected, it would no longer be present. We can then say that *T* is present *because* it produces *E*. Wright makes the point vividly:

If an organ has been naturally differentially selected-for by virtue of something it does, we can say that the reason the organ is there is that it does that something. Hence we can say animals have kidneys because they eliminate metabolic wastes from the bloodstream; porcupines have quills because they protect them from predatory enemies; plants have chlorophyll because chlorophyll enables plants to accomplish photosynthesis; the heart beats because its beating pumps blood (L. Wright 1973, 159).

Neander (1991a,b) provides a more rigorous formulation of this idea. Neander proposes to define the notion “proper function” as follows:

It is a/ the proper function of an item (X) of an organism (O) to do that which items of X 's type did to contribute to the inclusive fitness of O 's ancestors, and which caused the genotype, of which X is the phenotypic expression, to be selected by natural selection (1991a, 174).^{1,2}

Neander then claims that a teleological explanation cites an effect that is a “proper function”. In other words, a teleological explanation of cites an effect of the explanandum

¹ Millikan (1984, 1989a,b, 1993) has independently developed a similar theory. I focus on Neander's formulation because Neander explicitly aims to apply her notion of “function” to teleological explanation in biology.

² Neander tends to abbreviate this definition to the claim that biological functions are “effects *for which* traits were selected by natural selection” (1991a, 168). I avoid this here. The idea of a trait being “selected for its effect” sounds suspiciously teleological, and is easily confused with the similar-sounding causal notion of “selection for a trait” in a population. It is the latter, not the former, that is embodied in Sober's (1984) “selection for/selection of” distinction.

trait that contributed to the (inclusive) fitness of its past bearers, and in virtue of which the explanandum trait (or its genotype, if one takes this to be the unit of selection) was selected.

2.2 A digression on functions

Neander's approach to teleological explanation begins by defining a trait's "proper function" as the effect in virtue of which it was selected. A teleological explanation is then said to work by citing a trait's "proper function". But couldn't we say more concisely that a teleological explanation cites an effect in virtue of which the trait was selected? Why bring functions into it?

For Neander, the motivation seems to be an intention that the SE theory kill two birds with one stone. Neander aims not only to render teleological explanation respectable, but also to provide us with a conceptual analysis of what "function" means in biology. There are reasons to doubt that Neander succeeds in hitting the second bird. If the SE theory accurately captured the meaning of "function" as biologists use the term, one would expect biologists to be cautious in attributing a function to a trait by inferring from its *current* effects. We would expect them to defer to palaeontological evidence on such questions. Yet, as many authors have noted (Cf. Amundson & Lauder 1994, Wouters 2003, Lewens 2004, Weber 2005, Griffiths in press), biologists routinely ascribe "function" while possessing little knowledge of the evolutionary history of the feature with which they are concerned. Niko Tinbergen (1963) warns against conflating questions of function with questions of evolutionary history in the study of behaviour: in his iconic "four whys", these are separate questions.

But Neander's two birds are not inseparable. The success or failure of the SE account of teleological explanations does not turn on whether it also supplies, as a by-product, a conceptual analysis of function. It is irrelevant to the concerns of this paper whether Neander's notion of "proper function" captures the meaning of "function" in biology. One might take teleological explanations to be legitimate, and accept what the SE theory has to say concerning how they work, while rejecting Neander's purported conceptual analysis. In this sense the two issues can be teased apart.

2.3 A problem for SE

Let us return to teleological explanations. On the SE theory, to say "the heart beats because beating pumps blood" is to imply that beating hearts were once selected, that this selection explains why animals have beating hearts today, and that this selection occurred because beating hearts contributed to the fitness of their bearers by pumping blood. To say that, "birds have wings because wing enable flight" is to say that wings were once selected, that this explains why today's birds have wings, and that wings were selected because they contributed to fitness by enabling flight. It is an attractively simple, intuitively plausible view. One problem for the theory is that it fails to account for why richly teleological turns of phrase (like "the heart beats *in order to* pump blood", or "the *purpose* of the heartbeat is to pump blood") should strike us as appropriate. But this is not pressing, since the defender of the SE theory can reply that these are informal ways of citing a trait's proper function.

There is, however, a more serious problem for the theory. The SE theory commits us to the position that, whenever a teleological explanation like "hearts beat because beating pumps blood" is legitimate, that trait was once selected. A trait is "selected" when its frequency increases relative to *actual alternatives* in a population. Neander's definition of function

makes plain this commitment to the past occurrence of selection of the explanandum trait.

Brandon (2006) shows how this central commitment breaks down into four claims:

For instance, if the (SE) function of the red color of a flower is to attract pollinators, then it must be true that: (a) at one time in the history of the lineage in question there was variation in flower color, red being among the variants; (b) this variation was heritable (usually, but not always, this means that the variation has a genetic basis); (c) selection, in the form of pollinator discrimination, acted directly on flower color, not on some correlate of flower color, favoring red over alternative variants; and (d) this selection within the population genetic context of the lineage led to the form and frequency of red flowers we see in the descendent populations today (Brandon 2006, 268).

But the commitment of the SE theory to past competition between actual variants whenever a teleological explanation is appropriate turns out to be rather restrictive. As Robert Cummins (2002) argues, such a commitment implies that the SE theory is not able to cover teleological explanations such as “sparrows have wings for flying”, or “humans have eyes for seeing”. The SE theory would take the former explanation to imply that wings were once selected because wings enabled flight, and would take the latter to imply that eyes were once selected because eyes enabled sight. Yet:

To think of the modern eye or sparrow wing as itself selected is ... to conjure up a scenario in which there is a population of sightless primates or wingless songbirds into which is born a sighted or winged variation whose progeny take over the land or air. No one, of course, really believes anything like this. Yet

something very like this is implied by neo-teleology—by the idea that eyes are there because they enable sight and wings because they enable flight (Cummins 2002, 169).

Cummins's argument plays on a truism concerning how evolution occurs: complex traits do not appear all at once, with a sudden macromutation. They are produced by *cumulative* selection for small variations on existing structures (see e.g. Dawkins 1986, Sterelny & Griffiths 1999). They may well *never* have been selected themselves, simply because they have never faced competition from actual alternatives in the population. This claim becomes particularly plausible when we bring to mind complex and fundamental body parts like bones, skin and major organs. These are not traits for which there were actual competing alternatives. Consider the question of why the heart beats: the actual evolution of the beating heart can be traced to selection of some beneficial mutation enabling the contraction of a tube in chordates, in a population where no hearts was present and there was no blood to pump (Cf. Olson 2006). Successive beneficial mutations would have been selected over the millennia, one by one. At no stage in this story is there selection of beating hearts in a population where beating hearts, by virtue of pumping blood, out-competed an alternative trait, such as non-beating hearts. There was only ever selection of tiny variations on existing structures.

Neander (1995) is aware of the cumulative nature of adaptation, but she does not see the damaging consequences this observation carries for her (1991a,b) articulation of the SE theory. Complex traits are built by cumulative selection, but this does not imply that they have ever themselves been selected. Indeed, this will not be the case unless actual competing alternative traits have existed in the population. But if these traits have never been selected

then, *a fortiori*, they have never been selected in virtue of their advantageous effects. On the SE picture, complex traits that have never competed with actual alternatives cannot be given a teleological explanation, and of course, on Neander's definition of the term, cannot even have a proper function.

Cummins intends his argument both as a *reductio* of the SE definition of function and as a refutation of "neo-teleology", a term Cummins coins for the view that the theory of evolution allows legitimate teleological explanations. It is a mistake, Cummins argues, to think that "Paley questions"—which ask why a complex adaptation is present, and which were paradigmatic of the natural-theological tradition defended by William Paley (1802)—can be given scientifically respectable teleological answers in light of the theory of evolution. Only a long-winded historical explanation will do, detailing the many intermediaries and contingencies in the evolutionary history of the trait. At the very least, Cummins concludes, we need to shelve our more ambitious teleological explanations concerning heartbeats, eyes, wings and so forth. If we are to offer teleological explanations at all, it seems we can offer them only where the explanandum trait really has competed against actual alternatives. It is likely, then, that most if not all canonical examples of everyday teleological explanations of biological features (which tend to focus precisely on complex features like eyes, wings and hearts) are illegitimate.

2.4 Escape routes

Perhaps the defender of the SE approach to teleology need not be too disheartened by Cummins's objection. It cannot be wholly effective as a refutation of the SE position, since it allows that the SE theory *can* account for why citing a trait's effect can explain its current presence in those cases in which the trait in question *has* been selected at some point

in the course of evolution. And this criterion will plausibly be met *some* of the time when we want to give a teleological explanation of a trait's presence. Even rather complex traits may have at one time faced competition from actual variants.

Moreover, although evolutionary explanations appealing to the effects in virtue of which a feature was selected may not be able to explain why complex traits exist, they can presumably explain, at a more fine-grained level, why complex traits have certain properties rather than others. Recall Brandon's (2006) example of the red flower. They may never have been selection of *flowers* over an alternative, but selection of *red* flowers over other colours of flower can explain why the flowers are *red*. Likewise, even if there has never been selection of hearts over a competing alternative, selection of *efficient* hearts over inefficient hearts can explain why we have hearts that pump blood *efficiently* rather than inefficiently. Even if there has never been selection of skin over a competing alternative, selection of *pale* skin over dark skin in regions where low sunlight limits vitamin D synthesis may explain why Caucasians have *pale* skin rather than dark skin (Cf. Loomis 1967, Jablonski & Chaplin 2000). The SE theory supplies a sense in which, in such cases, it can be biologically respectable to say, for instance, "Caucasians have pale skin *because* pale skin increases vitamin D synthesis". Cummins shows that an explanation citing selection of the explanandum trait cannot plausibly be given for the presence of organs like hearts and skin, because there has never been selection of hearts in a population where some individuals lacked hearts, or selection of skin in a population where some individuals lacked skin. But perhaps the SE theorist can bite this bullet, and focus only on legitimizing teleological explanations of small-scale, selected features like pale skin and efficient hearts.

But I think there is a motive for exploring an alternative escape route for the shorthand approach, one that involves dodging this bullet rather than biting it. Cummins is right that when a complex trait has been built by cumulative selection of small variations, we cannot assume that the trait itself has competed against actual alternatives. Indeed, in many cases this possibility seems remote. But the moral I draw from this is not that teleological explanations of complex traits cannot be respectable. I think rather that it is precisely in virtue of this process of cumulative selection that teleological explanations *are* respectable. It is because hearts were built by cumulative selection that there is a sense in which they are there in order to pump blood. It is because wings were built by cumulative selection that there is a sense in which they are there to enable flight.

The essence of my proposal is simple. Teleological explanations point to the effects in virtue of which a trait is in some sense *optimal* for some organism in its environment. And traits that are optimal are the sort of traits we should expect a process of cumulative natural selection to throw up, provided a number of important assumptions obtain. This, in a nutshell, is why citing a trait's effects can explain its presence. The remainder of this paper presents this idea rather more thoroughly. In the next section, I outline the form of explanation at work here, which I term "landscape explanation". I then show how a landscape explanation lends itself to a teleological shorthand.

3. Explaining with Landscapes

3.1 Zooming out


We can appeal to selection to explain why the frequency of some trait changes over time, but this is not the only explanatory work the concept of selection can do. Another sort of

explanation appeals to selection to explain *why the trait came to exist in the first place*. Such explanations “zoom out”: they appeal to the *cumulative* effect of selection for beneficial variations on existing structures. This process builds complex traits gradually, through a series of intermediaries. One might suppose (as Cummins does) that answering a “Paley question” in this way requires detailed historical knowledge of the population in question, including fine-grained detail of the intermediaries that were selected. I contend that is not the case. Once one understands the basic process of cumulative selection, one need not give a longwinded story of fine-grained causal detail to answer “Paley questions”. There is a concise form of evolutionary explanation that does not require detailed knowledge of this sort.

This key, I suggest, is to visualize the cumulative effects of selection as the uphill movement of a population towards peaks on a *fitness landscape*. I contend that, from this zoomed out perspective, the *shape* of the landscape can explain why certain traits were the end products of cumulative selection. When the landscape is reasonably smooth and reasonably constant over time, we do not need fine-grained knowledge of the actual trajectory taken by the population across the landscape, since cumulative selection will push the population to the same peak from any of a wide range of initial conditions and trajectories. We do not need contingent historical detail of specific selection processes. We need only know the locations of the peaks. In this section, I spell out this picture in more detail. I first spell out the general form of explanation that I take to be at work here.

3.2 A simple case

Consider a ball dropped onto a landscape at time t (Figure 1). It rolls around and eventually comes to rest in the centre at t' . Why did it come to rest here? We can answer this question



with a causal story that cites the initial position of the ball, the gravitational force acting on it, and the subsequent trajectory that followed as a causal consequence from the initial position and the action of gravity. But suppose we do not have access to such information. Suppose we do not

know the ball's initial position, or the trajectory it subsequently took. All we know is that it was dropped somewhere on the surface, and ended up in the centre. Can we still explain its final position?

Even in this situation of relative ignorance, we can explain why the ball came to rest where it did. Because we know the shape of the landscape, we know that the centre is its lowest point, and so is the point at which the ball has least gravitational potential energy. We also know that the surface is reasonably smooth, so that a possible trajectory to the centre would have been available from almost any initial condition. On such a landscape, the ball will eventually settle at the centre almost regardless of its initial position and subsequent trajectory. It is not *quite* true that the ball will settle at the lowest point of the surface for *any* initial position, even on a reasonably smooth landscape like that shown in Figure 1. If the ball were placed right at one of the corners of this landscape, it would drop off the edge. But the existence of exceptional initial conditions does not invalidate our explanation when the ball is found to have settled in the centre. The “landscape explanation” works because a large range of initial conditions and subsequent trajectories will result in the ball ending up in the centre. Fine-grained causal knowledge of the ball's trajectory is unnecessary.

Explanations of this general form explain an observed state of a system (in this particular case, a particular position for the ball) by showing that this state has a special property: it is the state in which the system will be found in the long run, given any of some range of possible initial conditions and subsequent trajectories. The explanation does not aim to trace the particular initial state or trajectory actually taken by the ball (i.e. the line in Fig. 1). The possibility of a longer, separate explanation citing this fine-grained information thus remains open, though it would not necessarily provide the sort of understanding we seek. The landscape explanation may be the more helpful answer.

Do landscape explanations fit established models of scientific explanation? I do not think the legitimacy of the explanation in the simple ball-on-a-landscape case hangs on this question, but it will be helpful to settle it here. Models of explanation are now somewhat multifarious: landscape explanations fall within the scope of some, but not all. Because such explanations do not attempt to trace the physical, causal processes leading up to the final state of the system, they are not causal explanations on Wesley Salmon's (1984) "Causal Mechanical" model. For Salmon, explanations cite local, spatiotemporally continuous causal processes that produced the phenomenon we want to explain. Landscape explanations fail to supply information of this kind. Merely stating that the centre is the lowest point on the landscape in Fig. 1 does not give information about the causal processes by which the ball in fact came to rest in the centre. A causal-mechanical explanation of the ball's final position would cite the fine causal detail of the ball's trajectory.

Nevertheless, explanations of this sort have previously attracted attention in the philosophy of science, notably in Elliott Sober's (1983, 1984) discussion of *equilibrium explanation*, and Kim Sterelny's (1996) discussion of the difference between *actual sequence* and *robust*

process explanations. Both authors note that a causal explanation of some phenomenon need not cite information concerning the actual sequence of events by which it was produced. They observe that we can also explain by showing how the explanandum would have occurred given any of a range of initial conditions and causal trajectories. I suggest that “landscape explanations” are a special case of this general sort of explanation.

James Woodward (2003, 6) argues that Sober’s equilibrium explanations are in turn merely a variety of causal explanation. On Woodward’s view, causal explanations answer *what-if-things-had-been-different* questions: they tell us how interventions to given variables would have affected the explanandum. On Woodward’s view, the ball-on-a-landscape explanation qualifies as a form of causal explanation because it tells us something about how interventions would have affected the explanandum. It does not, however, show how the explanandum would have been *different* under interventions, but rather shows the opposite:

that intervening on the initial position or subsequent path of the ball within certain limits would have made *no* difference to the ball’s final resting place.



Fig. 2. Conformers of cyclohexane (lines represent carbon-carbon bonds, vertices represent carbon atoms).

3.3 Energy minima

The general notion of a landscape explanation can illuminate explanatory practices across the sciences. Before proceeding to biology, let me stress its importance in chemistry—a comparison between the two cases will prove fruitful later on. Take, for instance, the case of cyclohexane (C_6H_6). Cyclohexane can theoretically adopt a number of different conformational isomers (or *conformers*). Obvious examples include the planar

conformation (Fig. 2, top), the “chair” conformation (middle) and the “boat” conformation (bottom). But in a typical sample of cyclohexane at equilibrium, one conformation is overwhelmingly dominant: the chair. Other forms of the molecule can only be isolated with extreme difficulty; many, including the boat, cannot be isolated at all. Why is the chair form dominant?

In this case, the micro-level causal history of a given cyclohexane sample is inaccessible to us, so an explanation citing the initial conditions and subsequent interactions of the system of molecules is not an option. But we can still give an explanation. The explanation proceeds by showing that, of all possible conformations, the chair is the most stable. It minimizes torsional strain, and consequently is of particularly low energy (cf. Clayden *et al.* 2001). The laws of thermodynamics then dictate that this particularly stable conformer will dominate when a system of cyclohexane molecules is at equilibrium.

This explanation is similar in character to the ball-on-a-landscape explanation we considered above. One can picture the cyclohexane molecules exploring a multidimensional *potential energy surface*, with one dimension for each degree of freedom of the molecule, on which each point represents a possible conformation the molecule could adopt. The chair is then the *minimum* on the surface. The tendency of actual molecules to occupy this conformation in preference to, say, that corresponding to the boat conformation can be explained by appeal to the shape of the surface (see Atkins & Friedman 2005, Ch. 8, for an introduction to the notion of potential energy surface, and Leventis *et al.* 1997 regarding the conformational energy surface of cyclohexane).

3.4 Fitness peaks

How, then, might this form of explanation apply in biology? The sort of explanation at work in the cyclohexane example applies straightforwardly to the explanation of the conformation of larger molecules, including biological macromolecules such as proteins. We typically explain why a protein folds to adopt a particular conformation by citing the fact that the resultant structure corresponded to an energy minimum (eg. Bryngelson & Wolynes 1987, Pace *et al.* 1996, Leeson *et al.* 2000). But it seems that similar explanations are unlikely to be effective for complex, macroscopic traits. We cannot explain, for example, why mammals have eyes, simply by pointing out that completed eyes are particularly thermodynamically stable. This might explain why mammalian eyes don't fall apart, and it might help explain how mammalian eyes develop, but it does not explain why eyes evolved in the first place.³

There is, however, a form of explanation closely analogous to the explanations we have so far discussed that is supplied by the apparatus of *evolutionary* biology. Such explanations make use of Sewall Wright's (1932) notion of a "fitness landscape". An




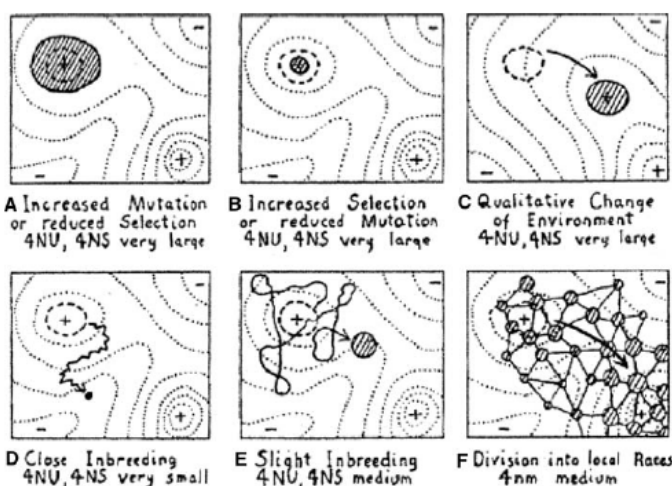
Fig. 3 from S. Wright (1932)

evolving population of organisms can be said to move around this notional landscape in a fashion roughly analogous to the ball on a real landscape. Note that the "fitness landscape", which often enjoys a central role in expositions of the fundamental ideas of the theory of evolution (e.g. Dennett 1995, Dawkins 1996, Godfrey-Smith 2009, Orr 2009), is distinct

³ Authors such as Brian Goodwin (1994) have sought an explanation for complex traits in complex systems theory rather than in evolution by natural selection, and may disagree with me here. In this paper I adopt the reasonably orthodox neo-Darwinian line that the existence of such traits is best explained by appeal to natural selection.

from the “adaptive landscape” familiar from textbooks (e.g. Futuyma 1998). On an adaptive landscape, mean population fitness is plotted as a continuous function of the frequency of alleles in a population. A population is then represented as a *dot* moving around the landscape as the alleles within the population change in frequency. By contrast, a *fitness* landscape represents the fitness of discrete individual *genotypes* (or, alternatively, phenotypes) in a given environment. Wright pictures “genotypes ... packed, side by side ... in such a way that each is surrounded by genotypes that differ by only one gene replacement” (S. Wright 1988, 116). The height of landscape at each discrete point shows the fitness of the genotype represented by that point. We can then represent a population as a *cloud* of genotypes, as shown in Wright’s original 1932 diagrams (Figures 3, 4). As

Fig. 4 from S. Wright (1932)



individuals appear, reproduce and die, the cloud moves around the landscape as the genotypic composition of the population changes.

What determines how close together two genotypes are found on the

fitness landscape? It is the similarity of the two genotypes, as measured by the number of genes that would have to be changed to get from one to the other. Individual genotypes that differ by relatively few genes will be closer together; genotypes that differ by many genes will be further apart. There is a complication here: one would need far more than three dimensions to represent these “distances” accurately. The problem with three is that, for a genotype G , the number of genotypes that differ from G by, say, 100 genes is vastly larger than the number that differ from G by ten genes, and this number is vastly larger than the

number that differ from G by exactly one gene. Yet a 3D landscape only allows the number of points at some radius from G to increase proportionally with the radius. A 3D fitness landscape cannot be a mathematically accurate representation of the “distances” between genotypes.

But a mathematical model was never Wright’s intention (see S. Wright 1988). The 3D fitness landscape is an explanatory tool that allows us to visualize some of the salient features of evolutionary dynamics on multidimensional landscapes. Importantly for our current purposes, it enables us to envision natural selection as a *hill climbing* process, tending to result in a population clustering around local peaks, just as a ball on a landscape tends to roll to the lowest point in its vicinity. Imagine a population on a smooth uphill slope leading to a peak. The fitter, “upslope” individuals have a greater expected number of offspring than their “downslope” competitors; and, since offspring tend to resemble their parents, this will tend to result in the relative number of individual genotypes in the higher region increasing over time. Overall, the population as a whole will tend to shuffle gradually uphill. When the random appearance of fitness-enhancing mutations results in genotypes even higher up the slope entering the population, selection will drive the population yet further uphill. The cumulative effect of the selection of new mutations will be the movement of the population towards a peak. When the peak is reached, no further uphill movement is available. New mutations will be selected against, and the genotypic composition of the population will be relatively stable.

The population may not reach the *global* peak: that is, the highest peak on the entire landscape for the environment in question. In sufficiently “rugged” landscapes, the population may get stuck on local peaks, separated from higher peaks by fitness “valleys”.

Wright (1932) makes much of this potential obstacle to adaptive evolution.⁴ Moreover, for various reasons, a population may not shuffle uphill at all. New mutations may fail to occur in the required direction, and in some cases frequency-dependent selection can cause a population to move *away* from a peak (see Moran 1964 for a number of examples where the “uphill push” fails to occur). A further complication (raised by Lewontin 1978) is that environments are unlikely to be wholly unchanging across any extended time interval, with the result that fitnesses of different genotypes—and thus the shape of the fitness landscape—may vary significantly over evolutionary time. Yet, despite these pitfalls, the general observation holds. Though populations will not always move uphill, the dynamics of natural selection suggest they typically will. When a range of assumptions hold, including the reasonable smoothness of the fitness landscape, the stability of its shape over a period of time, the non-frequency-dependence of selection and the regular appearance of new random mutants, a population will tend to move upwards until it reaches a fitness peak, at which point, if it gets there, it will be relatively adaptive and relatively evolutionarily stable.⁵

3.5 Explaining Evolved Features

Let us now consider how the notion of a fitness landscape might help explain the typical presence of a trait in a population of organisms. Suppose you ask why trait T is present in a population P . Why do humans have hearts? Why do sparrows have wings? We could give a fine-grained historical explanation of these phenomena. But I propose that we can also give another sort of explanation. In its general form, it proceeds as follows: *T is present because*

⁴ Kauffman & Levin (1987) further discuss the features of rugged fitness landscapes.

⁵ Note that a reliable mapping of genotype to phenotype in a given environment is required for the notion of a fitness landscape to work. One could avoid the need for this assumption by talking of fitness landscapes constituted by *phenotypes* rather than *genotypes*, but I do not explore this avenue here.

genotypes whose phenotypes contain T represented a peak on the fitness landscape of P or of an ancestral population from which P is descended.

Of course, for this explanation to work at all, T must actually correspond to a peak on a current or ancestral fitness landscape. The presence of a trait can only be explained by appeal to the shape of a given fitness landscape when it corresponds to a local optimum on that landscape. Moreover, certain assumptions must hold, viz. the assumptions which must obtain for hill-climbing to occur at all. These include the assumptions that the relevant landscape must be reasonably smooth and reasonably constant over time. The landscape must have been smooth enough for the peak corresponding to T to have been ascended from any of a wide range of initial positions. And the landscape must have been constant enough for the population to have had to time to ascend the peak corresponding to T from a similarly wide range of starting points.

In the ball example, the explanation of why the ball stopped where it did cites that the place it stopped was a potential energy minimum. In the cyclohexane example, the explanation of why one particular conformer is prevalent proceeded by showing that the conformer represented an energy minimum on an energy surface. Now, I have suggested that we can explain the presence of a trait in analogous fashion: by citing that genotypes with the actual trait represent a fitness peak on the fitness landscape of a current population or of its ancestral populations, provided a range of assumptions hold. Fitness-landscape explanations work in much the same way as the landscape explanations we considered in §3.2 and §3.3. They tell us that the explanandum trait would have evolved given any of a range of initial conditions and causal trajectories. But they do not presuppose the occurrence of past selection of the explanandum trait.

One may wish to supplement a landscape explanation with historical detail concerning the evolution of the trait in question. I do not claim that a landscape explanation is ever a *substitute* for a historical explanation of the trait's presence. But nor is a historical explanation a substitute for a landscape explanation. A landscape explanation offers understanding of a sort that a historical explanation may not be able to provide. It tells us that the evolution of a trait did *not* depend on the fine-grained details of a population's trajectory across the fitness landscape, but rather depended on that trait having the property of corresponding to a peak.

4. Teleology Revisited

Explanations that make use of a fitness landscape rely on a broad principle that populations tend to occupy certain regions of the fitness landscape *because those regions are peaks*. This idea suggests a way in which the effects by means of which a trait contributes to fitness can help explain why that trait is there. By considering the fitness-enhancing effects a trait would have, we can infer that genotypes with that trait are likely to constitute peaks on the fitness landscape, much as, for instance, considering torsional strain allows us to infer that the chair conformation of cyclohexane is at the minimum on that molecule's energy surface. Then, just as the relative energies of the conformations of cyclohexane explain why one is preferred, the relative heights of genotypes on the fitness landscape can explain why organisms in the population have evolved those genotypes through cumulative selection. There is thus a role for knowledge of a trait's effects in a landscape explanation of its presence.

Moreover, in contexts where the background theory and necessary assumptions are taken for granted, showing that the explanandum trait corresponded to a local optimum may be all the information we need to explain its presence. Talk of hill climbing on fitness landscapes can be elided. Consider the following example, supplied by Tim Lewens (2007):

Suppose a selection process is at work on some slow-running wolves. The wolves' environment may be such that *were* these wolves to run faster, they *would* catch more deer. ... It is thus legitimate to say that a particular pack of wolves is composed of fast runners *because* running fast helps wolves to catch deer (Lewens 2007, 55).

Taken in isolation, this observation seems puzzling. How can the fact that fast running *would* be advantageous for the wolves, if only they could do it, explain why they actually evolved to run fast? It sounds as though natural selection is an intentional designer, choosing which traits to build by contemplating their hypothetical advantageousness. But the notion of a fitness landscape can help make sense of such cases. Because running fast helps wolves to catch deer, and this advantage outweighs any potential drawbacks, the claim that fast running corresponded to a peak on the fitness landscape for these wolves is plausible. The "teleological" explanation is a concise means of drawing attention to this. We need not go into detail about the evolutionary trajectory by which the wolves have ended up fast. Presumably, change occurred via selection for small increases in speed, and hence it would be misleading to speak of fast-running wolves being selected over slow-running wolves. Nevertheless, ignorance of the actual course of selection does not preclude a landscape explanation.

On the view I am proposing, a landscape explanation can lend itself to a teleological shorthand that cites only the advantageous effects of a trait. To say that “sparrows have wings because wings enable flight” is respectable because it implies the following lengthier explanation:

Cumulative selection acting on sparrows or an ancestor of sparrows resulted in wings because wings corresponded to a peak on the relevant fitness landscape—because wings enabled flight.

The longhand explanation does not presuppose design, nor does it posit selection of winged birds over wingless birds in a historical population. Likewise, to say that “the human heart beats because beating circulates the blood” is respectable because it implies the following longhand:

Cumulative selection acting on humans or an ancestor of humans resulted in beating hearts because beating hearts corresponded to a peak on the relevant fitness landscape—because beating hearts circulate the blood.

To underline the point, this landscape explanation does *not* presuppose historical selection for beating hearts over non-beating hearts. The teleological shorthand is rather more concise than the evolutionary longhand. I suggest that there is no harm in using this shorthand, if we are clear about what is implicit.

5. Conclusion

5.1 The big picture

The “selected effects” approach to teleology gets something right: it is motivated by a realization that the process of natural selection works in such a way as to make teleological explanations respectable. This fundamental insight in part explains, I suspect, why the view has garnered such widespread support within both biology and philosophy. But, as Cummins (2002) notes, explicit formulations of the SE theory cannot cover teleological explanations of complex traits. SE theories focus on why a trait was selected, yet adaptation, generally speaking, occurs by *cumulative* selection. One should not imagine that complex traits (like wings, eyes and hearts) have achieved and maintained their current prevalence by out-competing actual alternatives. When we ask why birds have wings or why hearts beat, the answer cannot be that winged birds out-competed wingless birds, or that beating hearts out-competed non-beating hearts. Rather, complex traits result from selection for small variations on existing structures.

Though Cummins hopes to precipitate the demise of “neo-teleology”, I have drawn a different moral from his observation. Cumulative selection does not just throw up complex traits at random, irrespective of their contribution to fitness. On the contrary, when we “zoom out” to visualize evolution over extended periods of evolutionary time, we can envision cumulative selection as a process that tends to push populations uphill on a fitness landscape, the topography of which is determined by the population’s environment. With some assumptions about the smoothness and constancy of that landscape over the time required for cumulative selection, we can start to construct explanations for why certain traits have been built by the process. We can explain the evolution of a trait by pointing out that genotypes with that trait sit atop peaks on the fitness landscape. I have argued that this form of explanation lends itself to a teleological shorthand in which the relevant

assumptions are implicit. If a trait T , by virtue of its effect E , makes a contribution to fitness which is such that the population sits atop a peak on the fitness landscape, we can correctly give a shorthand explanation that simply says, “ T is there because it produces E .”

This account of teleological explanation is in some sense “etioloical”. Like the SE theory, it yields the result that teleological explanations in biology are legitimate, without any presupposition of backward causation or intentional design, because they imply an evolutionary explanation. But I have argued that the sort of evolutionary explanation that lends itself to a teleological shorthand in biology is more “zoomed out” than the SE theory suggests. It does not cite selection for the explanandum trait. It appeals instead to the “big picture”: the population’s fitness landscape.

I want to close by considering how the theory of teleological explanation sketched in this paper can provide tentative answers to two further questions regarding the applicability of teleological explanation in biology. By thinking of teleological explanations as abbreviated landscape explanations, we can, I think, suggest criteria for when a teleological explanation of a trait is biologically respectable, and suggest a reason why teleological language seems appropriate when we explain a trait’s presence by appeal to its effects.

5.2 When a teleological explanation is legitimate

Whenever philosophers seek to account for teleological explanations, the threat of a sceptical reply lurks in the background: *they are not genuine explanations at all!* After all, it is not obvious from their surface form that teleological explanations in biology are anything more than a relic from Paley’s natural theology, a careless consequence of our inclination to see organisms as if they were intentionally designed when they are not. There is plausibility

in the rebuttal that runs: “You can’t explain why something is *there* just by saying what it *does*! You’d have to be a Creationist to think such a thing!” Philosophers from Carl Hempel (1959) to Robert Cummins (2002) have sought to expose teleological explanations as shaky edifices built on false assumptions, while biologists from Michael Ghiselin (1974) to David Hanke (2004) have sought to expunge teleological explanations from their field (it is, in Hanke’s eyes, “the explanation that bedevils biology”). Such criticisms cannot be taken lightly. For those who seek to defend “neo-teleology” in biology, the project is to elucidate when and why citing a trait’s effect can help explain its presence. The goal is to make the sceptical reply look rather less plausible than it does at first glance.

I have aimed to meet this challenge. According to the “landscape” account, teleological explanations imply a form of explanation that appeals to the structure of a population’s fitness landscape. The model makes clear the situations in which teleological explanations of a trait are appropriate: they work when the conditions required for the longhand landscape explanation are met. “Hill climbing” must have occurred by cumulative selection of beneficial mutations, and the relevant fitness landscape on which the hill climbing occurred must have been reasonably constant and reasonably smooth, such that the peak could have been reached from any of a wide range of starting points and trajectories. Whether these conditions have been met in a particular case is an empirical matter. Landscape explanations demand a rich background of assumptions. If it turns out that these assumptions do not hold in a particular case, a teleological explanation will be spurious in that case. How often teleological explanations are legitimate in biology will depend on how often these assumptions hold.

5.3 Why teleological language strikes us as appropriate

I noted at the start of this paper that teleological explanations often do not simply state a trait's effects. They tend to make use of explicitly teleological terminology. Hearts beat *in order to* pump blood. Wings are *for* flying: flying is their *purpose*. I pointed out that we typically associate this sort of language with artefacts and agents. Defenders of a shorthand approach to teleological explanations aim to dispel a worry that arises from this, viz. the worry that teleological explanations presuppose intentional design. But this still leaves an unanswered question: if organisms are not designed artefacts, in virtue of what is teleological language appropriate? On John Maynard Smith's (1990) view, it is merely a matter of being concise. Phrases like "in order to" are all part of the effort to convey an evolutionary explanation as compactly as possible. But compare, "hearts beat in order to pump blood" with, "hearts beat because beating pumps blood". The first is no more concise for inserting a teleological "in order to". What tempts us to use such terms? I want to close by suggesting a speculative answer.

A ball dropped into a crater rolls around and tends to settle in the centre. Is it unreasonable to think of the ball as though it were "seeking out" the point of lowest potential energy? In this case, maybe it is a little far fetched. The ball is changing position, but not by virtue of any internal changes. But on the sort of energy surface we considered in the case of cyclohexane (Section 3.2), there is a difference: now the points on the landscape do not simply correspond to different energies a molecule could have without exhibiting any change in form, but rather correspond to *different forms* a molecule could adopt. Thus, trajectories of a system of molecules across the energy landscape can be visualized as molecular rearrangements. And as we saw in Section 3.2, cyclohexane molecules tend to arrange themselves in such a way as to ensure minimum energy and maximum stability. Now, I think, we are much closer to apparent intentional "seeking out" behaviour. It is as

though the molecules are *organizing themselves*. An evolving system of molecules settling on a preferred conformation behaves, to our eyes, as though striving towards the lowest energy state.

This conjecture entails a prediction: that, when chemists talk about molecules rearranging from a higher energy to a lower energy state, they should find it natural to use teleological language. And they do—particularly in pedagogical explanatory contexts. A standard undergraduate textbook, for example, tells us that "alkyl migrations occur in order to make a carbocation more stable" (Clayden *et al.* 2001, 983). Teleological language provides a basic heuristic framework for questions and explanations. For instance:

You now know that carbocations rearrange by alkyl shifts to get as stable as they can be—but this carbocation is already tertiary, and there is no ring strain, so why should it rearrange? (Clayden *et al.* 2001, 984).

One could cite many more examples. A recent empirical study of teleological explanations in chemistry teaching (Talanquer 2007) concludes that:

The occurrence of teleological explanations is tightly linked to the existence of a rule, principle, or law that governs the behaviour of the system, and that explicitly or implicitly implies the minimisation or maximisation of some intrinsic property (e.g., total energy, entropy, free energy). This law or principle tends to provide a sense of preferred direction in the evolution of a transformation (Talanquer 2007, 8).

In a nutshell, the position I have defended in this paper is that teleological explanation in biology works in much the same way. Because evolving populations tend to shuffle uphill on a fitness landscape, we can explain why they evolve certain features by pointing out that organisms with certain traits occupy peaks on the landscape. I want to end by speculating that, because evolving populations tend to move towards peaks, we are inclined to envision them as though evolving “in order to” reach a peak. Just as cations rearrange “in order to” get to the most stable state, we evolved eyes “in order to” be able to see, and birds evolved wings “in order to” to be able to fly. Landscape explanations naturally lend themselves to this form of shorthand.

References

- Allen, C., M. Bekoff & G. V. Lauder, eds. (1998), *Nature's Purposes*. Cambridge, MA: MIT Press.
- Amundson, R. & G. V. Lauder (1994), "Function without Purpose: The Uses of Causal Role Function in Evolutionary Biology", *Biology and Philosophy* 9: 443-469.
- Ariew, A., R. Cummins & M. Perlman, eds. (2002), *Functions: New Essays in the Philosophy of Psychology and Biology*. Oxford: Oxford University Press.
- Atkins, P. W. & R. S. Friedman (2005), *Molecular Quantum Mechanics*. 4th ed. New York: Oxford University Press.
- Ayala, F. J. (1970), "Teleological Explanations in Evolutionary Biology", *Philosophy of Science* 37: 1-15.
- Bechtel, W. (1989), "Functional Analyses and Their Justification", *Biology and Philosophy* 4: 159-162.
- Beckner, M. (1959), *The Biological Way of Thought*. New York: Columbia University Press.

- Brandon, R. (1981), "Biological Teleology: Questions and Explanations", *Studies in History and Philosophy of Science* 12: 91-105.
- — — (1990), *Adaptation and Environment*. Princeton, NJ: Princeton University Press.
- — — (2006), "Teleology in Self-Organizing Systems", in B. Feltz, M. Crommelinck & P. Goujon, eds., *Self-Organization and Emergence in Life Sciences*, Dordrecht: Springer Netherlands.
- Bryngelson, J. D. & P. G. Wolynes (1987), "Spin Glasses and the Statistical Mechanics of Protein Folding", *Proceedings of the National Academy of Sciences of the United States of America* 84: 7524-7528.
- Buller, D. J. (1998), "Etiological Theories of Function: A Geographical Survey", *Biology and Philosophy* 13: 505-527.
- — —, ed. (1999a), *Function, Selection and Design*. New York: SUNY Press.
- — — (1999b), "Natural Teleology", in (Buller 1999a), 1-27.
- Clayden, J., N. Greeves, S. Warren & P. Wothers (2001), *Organic Chemistry*. Oxford: Oxford University Press.
- Cummins, R. (2002), "Neo-Teleology", in (Ariew *et al.* 2002), 157-173.
- Dawkins, R. (1986), *The Blind Watchmaker*. New York: W. W. Norton & Company.
- — — (1996), *Climbing Mount Improbable*. New York: W. W. Norton & Company.
- Dennett, D. C. (1995), *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon & Schuster.
- Futuyma, D. (1998), *Evolutionary Biology*. 3rd ed. Sunderland, MA: Sinauer.
- Ghiselin, M. (1974), *The Economy of Nature and the Evolution of Sex*. Berkeley: University of California Press.
- — — (1994), "Darwin's Language May Seem Teleological, but His Thinking is Another Matter", *Biology and Philosophy* 9: 489-492.

- Godfrey-Smith, P. (2009), *Darwinian Populations and Natural Selection*. New York: Oxford University Press.
- Goodwin, B. (1994), *How the Leopard Changed Its Spots: The Evolution of Complexity*. London: Phoenix.
- Griffiths, P. E. (in press), "In What Sense Does 'Nothing in Biology Make Sense Except in the Light of Evolution'?", *Acta Biotheoretica*.
- Hanke, D. (2004), "Teleology: The Explanation that Bedevils Biology", in J. Cornwell, ed., *Explanations: Styles of Explanation in Science*. Oxford: Oxford University Press, 143-157.
- Hempel, C. G. (1959), "The Logic of Functional Analysis", in L. Gross, ed., *Symposium on Sociological Theory*. New York: Harper and Row, 271-287.
- Jablonski, N. G. & G. Chaplin (2000), "The Evolution of Human Skin and Skin Coloration", *Journal of Human Evolution* 39: 57-106.
- Kramer, P. J. (1984), "Misuse of the Term Strategy", *Bioscience* 34, 405.
- Kauffman, S. & S. Levin (1987), "Towards a General Theory of Adaptive Walks on Rugged Landscapes", *Journal of Theoretical Biology* 128: 11-45.
- Leeson, D. T., F. Gai, H. M. Rodriguez, L. M. Gregoret & R. B. Dyer (2000), "Protein Folding and Unfolding on a Complex Energy Landscape", *Proceedings of the National Academy of Sciences of the United States of America* 97: 2527-2532.
- Leventis, N., S. B. Hanna & C. Sotiriou-Leventis (1997), "A Three-Dimensional Energy Surface for the Conformational Inversion of Cyclohexane", *Journal of Chemical Education* 74: 813-814.
- Lewens, T. (2004), *Organisms and Artefacts: Design in Nature and Elsewhere*. Cambridge, MA: MIT Press.
- — — (2007), *Darwin*. London: Routledge.

- Lewontin, R. (1978), "Adaptation", *Scientific American* 239: 212-228.
- Loomis, W. F. (1967), "Skin-Pigment Regulation of Vitamin-D Biosynthesis in Man",
Science 157: 501-506.
- Maynard Smith, J. (1990), "Explanation in Biology", in D. Knowles, ed., *Explanation and its Limits*. Cambridge: Cambridge University Press, 65-73.
- Mayr, E. (1974), "Teleological and Teleonomic, A New Analysis", in R. S. Cohen and M. W. Wartofsky (eds.), *Boston Studies in the Philosophy of Science Volume XIV: Methodological and Historical Essays in the Natural Sciences*, Dordrecht: Reidel, 91-117.
- Millikan, R. G. (1984), *Language, Thought and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- — — (1989a), "An Ambiguity in the Notion 'Function'", *Biology and Philosophy* 4: 172-176.
- — — (1989b), "In Defence of Proper Functions", *Philosophy of Science* 56: 288-302.
- — — (1993), *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.
- Moran, P. A. P. (1964), "On the Non-Existence of Adaptive Topographies", *Annals of Human Genetics* 27: 383-393.
- Nagel, E. (1977), "Teleology Revisited", *Journal of Philosophy* 74: 261-301.
- Neander, K. (1991a), "Functions as Selected Effects: The Conceptual Analyst's Defence",
Philosophy of Science 58: 168-184.
- — — (1991b), "The Teleological Notion of 'Function'", *Australasian Journal of Philosophy* 69: 454-468.
- — — (1995), "Explaining Complex Adaptations: A Reply to Sober's 'Reply to Neander'",
British Journal for the Philosophy of Science 46: 583-587.

- Olsen, E. N. (2006), "Gene Regulatory Networks in the Evolution and Development of the Heart", *Science* 313: 1922-1927.
- Orr, H. A. (2009), "Fitness and Its Role in Evolutionary Genetics", *Nature Reviews Genetics* 10: 531-539.
- Pace C., B. Shirley, M. McNutt & K. Gajiwala, "Forces Contributing to the Conformational Stability of Proteins", *Journal of the Federation of American Societies for Experimental Biology* 10: 75-83.
- Paley, W. (1802), *Natural Theology or Evidences for the Existence and Attributes of the Deity Collected from the Appearances of Nature*. London: Faulkner.
- Pigliucci, M. (2008), "Sewall Wright's Adaptive Landscapes: 1932 vs. 1988", *Biology and Philosophy* 23: 591-603.
- Plutynski, A. (2008), "The Rise and Fall of the Adaptive Landscape?", *Biology and Philosophy* 23: 605-623.
- Salmon, W. C. (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- — — (1989), *Four Decades of Scientific Explanation*. Minneapolis, MN: University of Minnesota Press.
- Sober, E. (1983), "Equilibrium Explanation", *Philosophical Studies* 43: 201-210.
- — — (1984), *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Chicago, IL: University of Chicago Press.
- Sterelny, K. (1996), "Explanatory Pluralism in Evolutionary Biology", *Biology and Philosophy* 11: 193-214.
- Sterelny, K. and P. E. Griffiths (1999), *Sex and Death: An Introduction to Philosophy of Biology*. Chicago, IL: University of Chicago Press.
- Talanquer, V. (2007), "Explanations and Teleology in Chemistry Education", *International*

Journal of Science Education 29: 853-870.

Tinbergen, N. (1963), "On Aims and Methods in Ethology", *Zeitschrift fur Tierpsychologie* 20: 410-433.

van der Steen, W. J. (1971), "Hempel's View on Functional Explanation: Some Critical Comments," *Acta Biotheoretica* 20: 171-178.

Walsh, D. (1996), "Fitness and Function," *British Journal for the Philosophy of Science* 47: 553-574.

Weber, M. (2005), *Philosophy of Experimental Biology*. Cambridge: Cambridge University Press.

Wimsatt, W. (1972), "Teleology and the Logical Structure of Function Statements," *Studies in the History and Philosophy of Science* 3: 1-80.

Wouters, A. G. (2003), "Four Notions of Biological Function", *Studies in History and Philosophy of Biological and Biomedical Sciences* 34: 633-668.

Wright, L. (1973), "Functions", *Philosophical Review* 82: 139-168.

— — — (1976), *Teleological Explanations: An Etiological Analysis of Goals and Functions*. Berkeley, CA: University of California Press.

Wright, S. (1932), "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution", *Proceedings of the Sixth International Congress on Genetics*, 355-366.

— — — (1988), "Surfaces of Selective Value Revisited", *American Naturalist* 131: 115.