Natural Selection and Dennett's Stances*

Proposing a New Stance for Selection Processes

Christopher von Bülow†

February 2014

Abstract

I argue against Daniel Dennett's contention that evolution by natural selection ('Mother Nature') can only be fully understood from the intentional stance. I agree with Dennett that the physical stance does not enable us to see the real patterns in natural selection. Dennett's two other stances, however, are both already more sophisticated than necessary: on the intentional stance, the notion of 'Mother Nature's beliefs' doesn't do any work (she simply does what she 'desires' to do); and the mechanism of natural selection can be credited with 'design' merely in the minimal sense that it does fulfill its 'function'. Instead, I propose a simpler framework with just the degree of conceptual sophistication needed for understanding natural selection (and similar processes): the selection stance.

Evolution and Dennett's stances

Dennett's stances and functional complexity

Dennett claims that evolution by natural selection must be viewed from the intentional stance (e. g., Dennett 1987a, 314–21). To me this suggestion seems weird and feels wrong: I can't see anything in the process of natural selection that would merit being considered as an intentional system, that is, as a rational agent. I think Dennett's three stances – physical stance, design stance and intentional stance – are in that order appropriate for systems of increasing sophistication, increasing functional complexity. The mechanism and the process of evolution by natural selection, however, even though their products are indeed incredibly complex, cannot be of great functional complexity, because they are not themselves 'designed' (as opposed to artifacts, which are designed by people, and living beings, which are 'designed' by evolution). I suppose evolution does merit the application of a stance beyond the physical stance, but it should be a stance even 'below' the design stance. Thus I feel compelled to introduce the intermediate 'selection stance'.

Functional complexity, sophisticated structure

What do I mean by "functional complexity"? – An ordinary stone is very complex, if you consider all the details: all the molecules making it up, and their arrangement; it is much more complex than a crystal, which is very regular and homogeneous and thus rather simple. But the stone is complex in a boring sense. Its structure may have certain effects, but it doesn't *do* anything, it isn't especially *good* for anything – the stone's structure isn't sophisticated (it is the kind of structure that can easily, and does

*Former working title: "Mother Nature and the Intentional Stance".

†eMail: Christopher.von.Buelow@uni.kn; Website: www.uni.kn/FuF/Philo/Philosophie/philosophie/index.php?article_id=88. – I am grateful to Lars Dänzer and Jacob Rosenthal for helpful comments and discussions.

usually, arise by mere chance). An artifact or a living being, by contrast, is complex in an interesting sense; its structure is sophisticated because it is surprisingly well suited for achieving certain effects. My 'functional complexity' and 'sophisticated structure' are, I suppose, really the same as Dennett's 'design'; but I tried to find words that do not already suggest that the feature to be explained is due to some agent.

The physical stance

Ordinary physical objects or systems (ones that aren't living beings, or parts of them, or produced by them) are very low in functional complexity; their structure isn't very sophisticated.[1] Their behavior can be understood, explained, predicted, and often even manipulated, already from the *physical stance* – to the degree they can be understood, explained etc. at all: some physical phenomena (e.g., the weather or the precise behavior of fluids) can be predicted (not to mention manipulated) to some extent only by dint of vast information-processing efforts; however, even with regard to these phenomena the physical stance is still the most useful one.

Cause and effect

This stance, I gather, consists in applications of cause–effect thinking: "If this happens under circumstances so-and-so then that happens", "If I did this to degree so-and-so then that would happen in measure so-and-so", "If somehow this were achieved then that would be prevented." There is no notion of good or bad, of right or wrong applicable here, as there is in matters of design and of intentionality. Some objects or systems are better designed for a given purpose than others; and beliefs, desires, inferences and behaviors may be more, or less, adequate to circumstances, more or less rational, given the interests of the agent concerned. But as far as the physical stance is concerned, things and systems just do whatever they do; there is no criticizing their features or their performance, all there is to adjust are one's own expectations.

The design stance: function and optimality

The physical stance can be superseded in utility by the *design stance* once functionally complex objects or systems are considered. For some applications the design stance can even be practically indispensable, e.g., for using software programs. The design stance works by assuming that the object or system of interest is there *for doing* something, or in other words, that there is a *function* or *purpose* it has,[2] and – an idealizing and thus strictly speaking contrafactual assumption – that it is *designed optimally* for this function, given whatever constraints the designing process or agent was working under. In a further step, one can then attribute to it some specific function and try making inferences on that basis.

No designing agent presupposed

Here, a system's having a 'function' or 'purpose' is not intended as implying that it has been made or selected for that purpose by some agent. At least that is my understanding of the design stance: that it shouldn't already presuppose the intentional stance. (Is this already the point where I part company with Dennett? In that case, see my remark on page 14.) Rather, attributing a function merely means assuming that the system's features (and its 'behavior', if such there is), insofar as they are relevant for achieving the function, are well suited to that purpose, and to each other; and that the effect postulated as the system's function is (or once was) indeed brought about by it. Further implications, I think, are that the system fulfills its function in a certain *way* or according to a certain 'method' (to assume such a

[1] This is not necessary, just so overwhelmingly probable that it can be considered certain for all practical purposes whatsoever. The more sophisticated the structure, the less likely it is that things having that structure will arise by chance instead of thanks to a designing process or agent: an ordinary, unworked stone may happen to be a serviceable hand ax; but wherever we stumble upon a pocket watch or a jellyfish, we needn't take into account the possibility that it might have arisen by lightning striking randomly into unsorted heaps of dead matter.

[2] By being there for doing F, the system is in a sense *about* F. Having a function is thus of a kind with intentionality (cf. Dennett 1987a).

method is tantamount to attributing a more detailed function); and that it therefore also fulfills, in a suitable order, the corresponding subsidiary purposes.

Example: a hammer

One may think that the design stance doesn't allow us to predict anything of interest about, for instance, a simple tool like a hammer, because it doesn't have a functionally very complex physical structure and it doesn't have any behavior to speak of. But that is wrong: Assume you find a *part* of a hammer – say, in an archaeological dig – and recognize it as such (or hypothesize that it is such); then you can infer quite a few things about the missing parts, because they probably once fit together with the present part so as to form a *working* hammer. For example, if you find the head of the hammer then you can infer things about the handle; e. g., by inspecting the hole in the head, you can infer that the handle must have had a certain shape, because one end of it must have fit snugly inside the hole. In this manner a paleontologist draws inferences about the body of a prehistoric animal just by investigating one of its fossilized bones. She can do this only thanks to the design stance. The physical stance allows (almost) no such inferences: in general, the physical features of the present part entail neither the existence nor the physical features of other, missing, parts. The lone exception I can see is when there's a fracture surface; and all that can be inferred then is that here there once was *more of the same*.

Malfunction

When there is malfunction the design stance breaks down (stops being useful) partly or wholly, depending on the type of the malfunction. Then you have to go back to the physical stance, or at least to an application of the design stance where more details of the actual design, the actual way in which the main function is pursued, are taken into account, in particular the design's weaknesses.

The intentional stance: beliefs and desires

It becomes appropriate or even necessary to switch from the design stance to the *intentional stance* when a system has such a sophisticated design that it is flexible in how to fulfill its purpose. This means it can to some degree adapt to varying circumstances, that is, it can change the strategy by which it currently seeks to fulfill its function, depending on the obtaining situation. We can then call its function its (ultimate) *goal* or *desire*, and the varying circumstances, as registered by the system, and the means–end relations making up the system's strategy for fulfilling its function, its *beliefs*.

The intentional stance is not an outgrowth of the design stance

It is not *necessary* for adopting the intentional stance towards a system to start from a design-stance interpretation. It is just that I believe that this is where intentional systems (systems worth considering from the intentional stance; Dennett 1971) normally come from: they are very sophisticated functional systems (my term, as far as I know, for 'designed' systems). With typical intentional systems like animals, we normally, upon observing their activities, adopt the intentional stance automatically; considering them from the design stance, as having a function, is something that tends to happen mostly in biological and philosophical discourse. Such very sophisticated functional systems will often display behaviors so involved that their true function may be obscured, the forest invisible because of the trees. Thus it is not exactly obvious that living organisms are, in the end, vehicles for the propagation of their genes.

The rationality assumption

To adopt the intentional stance toward a system means considering it as a rational agent: we attribute to it particular beliefs and desires and assume that it is rational, and thence explain, predict and possibly manipulate what it does. Prediction may go by vicarious means–end reasoning: "He believes there is an obstacle in front of the car he is driving; he wants to avoid crashing his car; therefore he will stop." (Analogously, only looking backwards, for explanation.) For prediction and manipulation to work, the attributions must of course be well chosen; not just any attribution of beliefs and

desires will work. Also, no actual agent will be 100 percent rational; any physical system will at best approximate perfect rationality to some degree. The less rational the system behaves, the more the specifics of its beliefs must be taken into account ("He wanted to stop, but he mistook the clutch for the brake pedal"), or even the design features of its perception and reasoning apparatuses.

The intentional stance: when is it necessary, when gratuitous?

To the degree that prediction from the intentional stance works well *and* at much lower information-processing cost than prediction from another stance (even if somewhat less reliably and with less precision), to that extent we are dealing with a genuine intentional system, a genuine believer, desirer and actor, and not just applying the intentional stance gratuitously toward a system that doesn't really require its adoption. For example, we could adopt the intentional stance to predict what a lone oxygen atom will do in the company of hydrogen atoms: "It is very eager to bind with other atoms because it wants to complete its electron shell, so it forms a bond with the closest two hydrogen atoms." But even though it is sometimes easier for us to envisage the behavior of atoms by regarding their dispositions as desires, such an adoption of the intentional stance is not really called for, because the added cognitive leverage as compared with the physical stance is minimal: we do not miss any real patterns (Dennett 1991) if we stay with the physical stance. (Similar remarks could be made about the relation between the physical and the design stance.)

Why stances?

Metaphysical choices

A few words about the 'stance stance' (cf. Dretske 1988) in general: Introducing a 'stance' other than the physical stance is a way of dealing with phenomena that *prima facie* seem to require metaphysical posits that go beyond the purely physical – for instance, functions, to explain teleological phenomena, or beliefs and desires, to explain intentionality. For a physicalist there are two obvious ways of treating such phenomena: *embracing* them by identifying them with certain physical phenomena, that is, being (strongly) realist about them, or *denying* them and thus being eliminativist about them. (One can argue about whether [versions of] reductionism and instrumentalism are realist or eliminativist; but I won't pretend to penetrate all the subtleties this topic involves.)

The stance stance

Dennett has found himself a place *between* these two stools, a position at least similar to instrumentalism which is sometimes called "interpretationism" (see also Dennett 1988, 536–8) and which he himself categorizes as "mild realism" (Dennett 1991, 98): Certain patterns are seen only when one attempts to interpret (understand etc.) a system by putting on the glasses of a certain stance (a theory?), that is, by making certain idealizing assumptions about the system (e.g., optimal design or rationality), which may be realized only imperfectly, and ascribing to the system appropriate properties (theoretical states?) of certain kinds (e.g., a certain function or certain beliefs and desires), and then following through what the rules and axioms (?) of the stance allow one to infer. If this strategy works reasonably well then one can predict the system's behavior much, much easier (if with less precision and reliability) than one could from the physical stance. In this case, one has found a 'real pattern': a pattern realized (approximately) in physical stuff but not itself a physical pattern, as reductionism or strong realism would have it.

Holism and indeterminacy

The difference to the postulation of theoretical entities and properties by ordinary, e.g., physical, theories seems to be that for the phenomena or patterns visible from a non-physical stance, *holism* obtains (see, e.g., Dennett 1981a, 58; 1987a, 290–93), that is, one can ascribe 'theoretical' properties only in whole packages, not singly,

independently from each other, and only to the system as embedded in a given environment, not in isolation. This holism entails a kind of *indeterminacy* (see, e.g., Dennett 1987b, 103–5), an occasional absence of facts of the matter about which ascriptions are correct, due to there being incompatible best-possible interpretations of the system that work equally well. But these are again questions I do not understand at all well. For the purposes of this essay, anyway, I accept all these tenets, except for Dennett's position about evolution having to be considered from the intentional stance.

Perspectives on natural selection

Ways of considering Mother Nature's work: global vs. local

Is there reason to adopt the intentional stance toward evolution by natural selection, a.k.a. Mother Nature? Note first that there are different ways this might be done: more local and more global ones, so to speak. The most *global* way of reflecting on natural selection would consider what effects it has on the whole of life on earth,[3] a more *local* way would focus on its effects on a single population, a single gene pool. Intermediate perspectives would look at whole species, genera and so on. Furthermore, one can also – reducing the scope still further – consider the effects of natural selection on a single gene and its alleles, or on larger parts of a species' genome.

Better not too global

On a 'global' point of view, adopting the intentional – *or* the design – stance seems clearly inappropriate: What goal (or function) could be attributed to evolution that would make it rational to, on the one hand, improve a predator species' proficiency at catching its prey, while, on the other hand, improving its prey's skill at evading capture? This kind of arms race is ubiquitous, and would suggest that Mother Nature is schizophrenic. Probably one *can* think up goals for her that make sense of such behavior (maybe she just likes functional complexity), but I suppose they would be rather unconvincing ones, and not really helpful for understanding evolution in general.

Local: where does the 'design' in a population come from?

We have a better chance of usefully employing the intentional stance toward evolution if we consider a single species or population and ask why it is as it is: Where did the high functional complexity shared by its members come from? How did this specific 'design' arise? That is, we ask the kind of questions that made the Reverend Paley infer the existence of a watchmaker from the finding of an isolated watch, the kind of questions the theory of evolution by natural selection is most urgently needed for answering. I suppose that this is also the perspective Dennett suggests we take.

Natural selection as an intentional system?

Rationales 'appreciated by Mother Nature'

So let's assume for the sake of argument that Mother Nature can be profitably viewed as an intentional system, and try at this level to find good attributions of beliefs and desires for her. The most frequent type of attribution I found in Dennett's writings (e.g., his 1987a) is that certain *rationales* or *reasons* are 'appreciated' or 'recognized' or 'acknowledged' (usually in scare quotes) by natural selection. For example, Dennett (1983, 258–9) writes about distraction displays by ground-nesting birds which seem intended to deceive predators. There is a complicated rationale for this kind of behavior:

[3] Of course there may be life on other planets or even elsewhere, but I'll stay with the one biosphere we can so far observe.

> I'm a low-nesting bird, whose chicks are not protectable against a predator who discovers them. This approaching predator can be *expected* soon to discover them unless I distract it; it could be distracted by its *desire* to catch and eat me, but only if it *thought* there was a *reasonable* chance of its actually catching me (it's no dummy); it would contract just that *belief* if I *gave it evidence that* I couldn't fly anymore; I could do that by feigning a broken wing, etc. (Dennett 1983, 258)

This piece of practical reasoning is written from the bird's perspective, but nothing like that could really have been entertained by *it*. Nevertheless the rationale is there:

> It is too obvious that the raison d'être of this instinctual behavior is its deceptive power. That's why it evolved. If we want to know why this strange dance came to be provokable on just these occasions, its power to deceive predators will have to be distilled from all the myriad of other facts, known and unknown and unknowable, in the long ancestry of the species. But who *appreciated* this power, wo *recognized* this rationale, if not the bird or its individual ancestors? Who else but Mother Nature herself? That is to say: nobody. Evolution by natural selection "chose" this design for this "reason." (Dennett 1983, 259)

Is Mother Nature a true believer?

So, the deceiving bird does not know what deception is and how it works, but Mother Nature does? Dennett writes: "Mother Nature ... That is to say: nobody"; so, is Mother Nature an agent or a nobody? Is she a 'true believer', "a system whose behavior is reliably and voluminously predictable via the intentional strategy" (Dennett 1981b, 15)? This may be what it means if we *have to* consider her from the intentional stance, if adopting it is more than just convenient but rather enables us to discern real patterns no other stance allows us to see. But attributing such an opinion to Dennett may be going too far. Probably he only claims that even though we do not have to classify natural selection as a true believer, a genuine rational agent, viewing it from the intentional stance is still necessary for understanding what goes on in evolution, for seeing those real, but nonphysical, patterns. I will however argue that this weaker claim is false too.

Mother Nature is a bit like a human breeder selecting for reproductive prowess

What beliefs, desires and actions should we attribute to Mother Nature? Thanks to Darwin, we already have a useful model for how to consider natural selection as an agent: It is somewhat like a farmer practicing *artificial* selection (selective breeding) on a type of animal or plant in order for it to obtain or increase certain traits he wants it to have. In the case of natural selection those organisms that have the 'desired' trait more strongly are 'allowed' by Mother Nature to reproduce more successfully, on average, than others. The trait 'desired' is of course precisely the ability to reproduce more successfully than others.[4] This, then, is what Mother Nature 'does': she 'selects' for comparative reproductive prowess (where the comparison is with the other members of the population); high reproductive prowess is what she 'wants' in her children. I believe that this is (the core of) the received view among biologists, just rendered in anthropomorphic (intentional) terms, and that Dennett would agree thus far.

[4]Or so goes the simplified story. In the more sophisticated story, brilliantly popularized by Richard Dawkins (1976), the variation Mother Nature selects from does not consist of (types of) organisms in a population but of the alleles of genes in some gene pool, and the quality she selects *for* is not reproductive ability but rather ability at self-replication via phenotypic effects. Yet the difference between these two stories – the organism- and the replicator-centered view of evolution – is irrelevant with respect to Mother Nature's intentionality (my arguments should carry over, *mutatis mutandis*), and so I will in the main stick with the more familiar organismic perspective.

Neither of the two stories is tautological: Relative aptness for successful reproduction (replication) does not guarantee *actual* successful reproduction (replication), it only statistically tends to give rise to it, the more so, the larger the number of apt organisms (carriers of tokens of the apt allele) involved. For large numbers, this probability becomes near-certainty.

Ways of life and derivative reasons for selection

Now, good reproduction as the key quality can necessitate countless subqualities. In order to reproduce, you have to live long enough, so you have to *make* a living in some way; and these ways are innumerable. Very roughly, you can be a plant or an animal; in the latter case you can be a herbivore, a carnivore or an omnivore; and so on, specializing ever farther, for different habitats and niches and circumstances. Each species has embraced one such intricately detailed 'way of life'; each population goes on further molding the way handed down to it. In this process of spelling out the details of a way of life one arrives at traits and behaviors whose connection to reproduction may be not at all obvious. Still, insofar as these qualities are ultimately conducive to reproductive success, Mother Nature will 'select' for them, too. She 'wants' these qualities in the members of the population *because* they enhance reproductive prowess, however mediately. Whatever else they may be good for, if they don't aid reproduction Mother Nature doesn't 'care' about them.

Mother Nature's beliefs and desires

I suppose these functional interrelations between different qualities might be considered as Mother Nature's *beliefs*: Say she 'believes' that (in population P) good eyesight aids survival; she further 'believes' that survival aids reproduction;[5] and consequently, since she 'wants' the members of P to reproduce as successfully as possible, she 'wants' them to be good at surviving, and therefore, to be good at seeing. For another population P′, say, one inhabiting dark caves, she may 'want' its members to have only rudimentary eyes, or none at all, because she 'believes' that in the case of P′ the resources needed to build and maintain sophisticated eyes are best invested in certain other purposes, which serve P′ better in survival and, ultimately, reproduction than good eyesight. Here we have Dennett's (1983, 259) 'free-floating rationales', considered as appreciated by Mother Nature.

Selecting physical options for functional reasons

Of course Mother Nature doesn't, in a given population, 'select' good eyesight *per se*; good or bad anything are not on the menu of physical options for her to 'choose' from. What she really 'selects' are genes whose usual phenotypic effects include better eyesight as compared to the effects of their alleles in the gene pool. And she does this over and over again, not only with successive generations of organisms but also with different genes and shifting collections of alleles for a given gene, building cumulatively on the 'design' that is already available. Only thus can she gradually 'climb Mount Improbable' (Dawkins 1996) to reach dizzying heights of functional complexity.

Intentional stance: appropriate if error is impossible?

Can Mother Nature 'err', can her 'beliefs' be false? If not, that would seem to indicate that the notion of belief doesn't do any work here; that the role supposedly played by Mother Nature's 'beliefs' is really played simply by the corresponding facts. This in turn would support my thesis that we are not dealing here with an intentional system in any interesting sense, and that something other than the intentional stance suffices.

How Mother Nature believes

Whether Mother Nature can 'err' depends on what her 'believing' consists in. How would she 'believe' truly, and how falsely, that good eyesight aids survival, or that it aids reproduction? Presumably she would 'believe' in these interrelations by 'letting' those members of the population that have better eyesight survive longer or reproduce more successfully, respectively, on average.

[5] More precisely, Mother Nature 'believes' that P-organisms with *better* eyesight will survive *longer*, and that those which live longer will reproduce more successfully – natural selection is always comparative. Also, we must ascribe to her some slightly less interesting further beliefs, namely, about *which* P-organisms have better eyesight (and should thus be 'allowed' to live longer) and which actually *do* live longer (and should thus be 'allowed' to reproduce more successfully).

How Mother Nature could err

And how could such a 'belief' be false? By coming about through mere accident. That is to say, it might happen that those members of the population which are visually deficient nevertheless live longer or have more offspring just because they have a lot of lucky strikes (e.g., in finding food or mates, or in avoiding predation), or because the other members are, for example, accidentally hit by lightning or rendered infertile by random cosmic rays.

Error possible on a tiny scale, as good as impossible in the long run

Something like this happening is plausible in the case where there are only a *few* keen-eyed individuals, and it is (just) conceivable in the case of hundreds, but it beggars belief that it should happen to great numbers of individuals over many generations. That an advantageous variant should in the long run (which is the only run that matters in evolution) accidentally *not* gain ground in the gene pool is of course in principle possible, but it is so astronomically unlikely that, were its semblance to be observed, we would start looking for the hidden deleterious side effects of the supposedly beneficial allele: does it somehow *attract* lightning? So, apparently Mother Nature can have false 'hunches', as it were, but we can safely trust that her stable 'beliefs' about what aids successful reproduction, the ones manifested in actual 'selection' (i.e., the ones solidified by fixation in gene pools), are all true.

Notion of belief: redundant

But then we'll never get into a situation where we have to explain the results of natural selection by saying, "Mother Nature chose quality Q_1 over Q_2 because she *thought* (erroneously) that Q_1 would allow P-organisms to reproduce better than Q_2 would." Rather it will always be, "She chose Q_1 over Q_2 because for the successful reproduction of P-organisms, Q_1 *is* better than Q_2." Mother Nature just 'selects' what is 'good'; that is, she simply does what she 'desires' to do. The notion of belief is not really needed, and that means we don't really need the intentional stance. Therefore I suggest that we look whether we can't understand evolution by natural selection from a stance that is 'lower' than the intentional one.

Applying the design stance to natural selection?

The design stance: evolution as machine-like

If we stick with Mother Nature's 'desires' (concerning a given population), we can switch to the design stance and try interpreting them as the functions (ultimate and derived) of a selecting apparatus. The primary 'function' of this selection machine would again be to improve reproduction in the population it operates upon, and its secondary, tertiary etc. improvement-'functions' would, as before, depend on the specific way of life of the population. This system would 'choose' certain physical (namely, genetic) features for spreading in the population, and physical qualities for enhancement, but again, it would not 'choose' them *for* their physical natures (visible from the physical stance), but for their *functional* or *adaptive* properties, namely, for the degrees to which they promote reproduction and its subsidiary activities – for properties invisible from the physical stance but visible from the design stance.

Natural selection works, but it is not well designed

This way of looking at natural selection, however, doesn't seem justified either, because natural selection is far from optimally 'designed' as an adaptation-producing mechanism – it is not 'designed' at all. For example, I do not see how one could make design-stance inferences from features of one part of this 'mechanism' to features of another (as we did in the example of the hammer). Admittedly, natural selection on replicators is vastly superior to raw, blind chance (where complex structures can only come about through one-shot attempts at randomly throwing together atoms, and are more likely to fall apart again than to be preserved and built upon). And needless to say, natural selection, if given lots and lots of time and toys to play around with, does produce exceedingly sophisticated 'design'.

Shortcomings of natural selection

However, comparing natural selection – unfairly, of course – to deliberate design by intelligent agents, at least three points one could complain about come to mind:

- its inefficiency: it is glacially slow and terribly wasteful;[6] even the stupidest options are tried out with live material;
- the undirectedness of mutations: the genetic variation that the 'evolutionary algorithm' (Dennett 1995) must work with is in no way prefiltered;
- the lack of foresight: change always has to come about by gradual improvement; nature cannot 'tolerate' a temporary decline in quality in order to reach new heights, and adaptive 'jumps' in design quality are virtually impossible.

Notion of function: vacuous

All the 'design' there is in the mechanism of natural selection – all the basis for adopting the design stance – consists in the fact that the function ascribed to it is actually fulfilled. But if "design" were to be taken in this minimal sense then we could ascribe functions to anything and everything: whatever regularly has an effect E could be considered as being there *for* doing E. The concept of function would have lost its substance.

A stance below the design stance, but above the physical stance?

So, while both the intentional and the design stance do capture patterns in natural selection that overtax the physical stance, it seems to me that they are methodological overkill. How can we make do with weaker assumptions and less demanding concepts without thereby losing sight of the relevant patterns altogether?

The selection stance

The selection stance: selecting for a quality

I suggest we give up both the idea of a selecting *agent* like Mother Nature and that of a selecting *apparatus*, a mechanism optimized towards enhancing some quality, and keep just the idea of considering a system or process as 'improving' its objects with respect to some quality by persistently, and perhaps cumulatively, 'selecting' *for* (higher degrees of) that quality and whatever qualities further it. I call this the *selection stance*. The idealizing assumption here is that, at least in the long run, 'progress' is actually made.

Target quality: good in the sense of selected for

The word "quality" is ambiguous; it can designate an arbitrary property or a good property. Here, it does a little of both: any property whatsoever may be selected for; but *insofar* as it is selected for it is 'good' relative to the selecting process, it is 'good' in the sense of being the ground on which things are 'chosen'. Matters will be more interesting if we add the restriction that this 'target quality' be a *graded* property, one that has varying degrees; otherwise the selection process becomes a very boring affair, finished as soon as all the 'have-nots' have been replaced by 'haves'.

Selection without agency or design

The term "selection", too, is used only in a weak sense, not implying conscious, deliberate choice but rather just a filtering process which can be implemented by wholly mindless as well as by more intelligent means: an ongoing reduction of the diversity in a set of tokens of various types which engenders a gradual growing of the proportion of tokens of 'better' types, finally leaving behind only (or almost only) tokens of the 'best' type(s), or of very 'good' types. *If* the diversity in the set of tokens (the population or the gene pool, in the case of evolution) is continually renewed by slight modifications of the remaining types (again quite possibly by mindless means),

[6] To be honest, though, I am not sure whether applying concepts like "wasteful" or "inefficient" to natural selection really makes sense, as alternative processes which might fulfill the same 'function' more efficiently seem bound to be completely different in character.

such a cumulative filtering process may be expected to converge upon sets of tokens of extremely 'good' types indeed. But as yet I tend to think that the assumption of renewed diversity should not be part of the selection stance but should rather constitute another optional extra.

Selection, selection for, and evaluation

What is selected *for* is the target quality, or better: increase in, or higher levels of, that quality; what is *selected* are the ('good') *types*. What is it that is being done to the individual tokens? Single tokens are not 'selected'; they just occur, or exist. As a term of art, "evaluation" seems to be fitting: the tokens are *evaluated* with respect to the target quality.

A precise rendering of the selection stance?

When we consider a particular application of the selection stance, that is, a specific set of tokens and a specific target quality we hypothesize is being selected for, then it should with luck be fairly clear which other qualities heighten the target quality and which lessen it, and to what degree. However, when we try to flesh out the selection stance in the abstract, when we try to give an abstract characterization of how to arrive, starting from an arbitrary set of initial conditions, at a prediction, then things get difficult. In order to make my hazy description of the selection stance somewhat precise, it would be proper to deliver at least the rudiments of a formal framework for it. My attempts at doing so have however not advanced very far yet and will need more time, effort and outside input.

The selection stance applied to evolution

Selection for comparative reproductive prowess

We can readily see how the evolution of a population of organisms by dint of natural selection constitutes an application of the selection stance: the tokens are the organisms in the population (or the bearers of various alleles of a given gene in the corresponding gene pool); the more or less 'good' types they instantiate are their traits and behaviors (or the phenotypic effects of different alleles); and the quality with respect to which 'selection' happens is reproductive (or self-replicative) prowess.

Derivative goodness

There are *derivative* forms of 'goodness', which consist in aiding reproduction (replication) in one way or another; tokens of such derivatively 'good' types are 'positively evaluated' not for any specific *physical* characteristic they possess,[7] but solely for being better at accomplishing something that promotes the ultimate 'good' of reproduction (replication). For instance, going back to Dennett's example of the bird's (call it B) distraction display cited on page 6, some of the relevant interrelations between qualities (abilities, in this case) are the following:[8]

- If B's chicks die, e.g., because they are eaten by a predator, that diminishes B's reproductive success; therefore if B can avert that danger it will (in comparison) boost its reproductive success.
- When a predator P is close to the nest containing B's chicks and thus constitutes a danger to their survival, then if B can lure P away from the nest, that will help avert the danger.
- If under these circumstances B can make P perceive B itself as easy prey then B can lure P away.

[7] However, I suppose the quality 'selected' for in applications of the selection stance may also be a purely *physical* property, in which case both primary and derivative 'goodnesses' would presumably be visible already from the physical stance.

[8] I omit the required "on average", "in general", "if cost permits" and so on. Also, what selection is really about is of course not whether organisms in the population are categorically 'able' or categorically 'unable' to achieve this or that, but rather differences in how *good* they are at achieving it.

- If B can feign a broken wing and attract P's attention, B can make P perceive it as easy prey.

Such interrelations between (design- and intentional-stance) qualities are Dennett's 'free-floating rationales': some traits just *are* good for reproduction (in a given population), and this has the effect of their being selected for as derivatively 'good' – without anyone's having *beliefs* about the connection.

Higher stances are needed – but not for Mother Nature

However, in understanding evolution by natural selection the envisaged network of functional interrelations between different qualities ("Which trait or ability promotes which, and to what degree?") is crucial. And it *cannot* be seen from the selection stance (nor the physical stance); it is a design- and sometimes even an intentional-stance affair. So, does the selection stance not suffice after all? Is the design (and the intentional) stance still needed for understanding evolution? I'm afraid this is indeed the case. But we do *not* apply the design or the intentional stance toward the evolutionary algorithm, toward Mother Nature herself; rather we do so only with regard to the question of what's good for what: "For what reason – for what other quality's sake – is this quality being selected for?" To answer this question, we adopt the design or the intentional stance toward different possible types of *organisms* and compare them to each other in how good they may be expected to be at achieving this or that and, in the end, reproductive success. In other applications of the selection stance, say, ones aimed at purely physical or geometrical qualities, the relations between the ultimate 'good' and the derivative ones may be visible from 'lower' stances.

Prediction

From the selection stance we predict that, in the long run, very 'good' types will be realized, that is, types of organism whose 'design' is sophisticatedly suited for successful reproduction in their given niche. If we know in some detail which types or qualities enhance reproductive prowess, directly or indirectly, then we can make more specific predictions. Will these leave out any predictions that one could make from the intentional stance? Are there any real patterns in evolution that *cannot* be seen from the selection stance – with a little help from the design or even the intentional stance – but can be, from the intentional stance? I submit that the answer is No.

From selection to design

How does the selection stance undergird the design stance in biological evolution; how does 'design' emerge from natural selection? – The quality selected for is a greater ability to *do* something. Cumulative selection then produces very great ability at doing that something, and such ability requires very sophisticated structure, high functional complexity – in other words, design (in the sense of the design stance).

Other applications of the selection stance

Genetic algorithms and memetic evolution

It may seem like the introduction of an additional stance for understanding evolution is completely *ad hoc*. But there are other candidates for application of the selection stance. It is obvious, for example, that this stance can be applied not only to biological evolution but to *all* cases where there are replicators, in particular, to the method of design via *genetic algorithms* (where the quality selected for is aptitude for solving a particular [type of] problem) and to natural selection among *memes* (which are 'evaluated' with regard to their capability for getting themselves imitated).

Neuronal group selection

In the latter case, however, I doubt that the selection-stance story of differential replication via imitation (in a wide sense) is already the whole story about cultural evolution: too much hangs on creative processes inside individual minds (von Bülow 2013, 8–9). The differential-imitation aspect of memetic evolution would have to be supplemented by an account of individual creativity – possibly due to a kind of natu-

ral selection in the brain. Such an account may be constituted by Gerald Edelman's theory of *neuronal group selection*, another obvious candidate for the selection stance. However, I know next to nothing about neurobiology, so I can only mention this as a suggestion to those who find the idea of the selection stance potentially worthwhile, and know more. It may however be necessary here to switch from the selection stance to the design stance when we are less interested in the (abstract, nonbiochemical) how of these processes and more in their why and their functions, because brains are evolved, that is, 'designed', organs in animal bodies.

Selection in immune systems

Another good candidate for adoption of the selection stance seems to be the functioning of *adaptive immune systems* in vertebrates, where, I understand, certain types of randomly generated antibodies are 'selected' according to the intensity of their defensive reaction to different pathogens. (I don't know whether we are dealing with replicators here, that is, whether the 'good' antibodies get multiplied with slight concurrent modifications and are then subjected to 'selection' again.) This is once more an example where, depending on the goals of inquiry, the design stance may be more appropriate than the selection stance.

Operant conditioning as selection

Then there is *operant conditioning*, where behaviors of an organism are 'selected' by the environment. Roughly, if a behavior token brings about a good (for the organism) event it is thereby 'rewarded', that is, the corresponding behavior type is reinforced; a behavior token with bad consequences is 'punished', its type is inhibited. Here, "reinforcement" and "inhibition" mean that these types become more or less likely, respectively, to be instantiated in the future. In this manner the repertoire of behavioral dispositions of the organism is honed towards adaptedness to the given environment (or that at least is the mechanism's purpose). The quality being 'selected' for in operant conditioning is a behavior's tendency to cause events that are good for the organism (by its own lights, anyway), rendering that type of behavior 'good' with regard to selection. I am not sure whether there are *derivatively* 'good' behaviors in operant conditioning – maybe when the organism has some understanding of which deeper dispositions (which cognitive 'behaviors') bring about its overt behavior.

Entropy increase as selection

A rather boring candidate for the selection stance (though not as boring as the sieving of sand) is the phenomenon of *entropy* in physics. I know very little about statistical mechanics, but I imagine that the increase of entropy in a closed system can be viewed from the selection stance by considering small configurations of particles as the tokens (and, in the abstract, as the types), and considering likelihood or stability as the quality 'selected' for, on the basis of simple probability.[9] Possibly the physical stance suffices here; but perhaps the Second Law of Thermodynamics (according to which the entropy of closed systems does not decrease), although part of physics, is not itself a straightforward law of nature, in being based on the idealizing assumption that certain extremely unlikely configurations (or trajectories through phase space) are never realized. So, even though the results will be unsurprising, the possible need for adopting the selection stance to correctly categorize entropy would seem interesting in itself.

Stance or theory?

Idealization, interpretation, holism, indeterminacy?

Now, is the selection stance as characterized here really a Dennett-style *stance* ("The Stance that Dan Forgot") or rather (more scientific, I guess, but philosophically less

[9] I suppose entropy increase is more or less what Richard Dawkins (1976, 13) means by "survival of the stable".

sexy) a framework for a very abstract *theory* of selection processes? Well, it employs an idealizing assumption which isn't necessarily true, namely, that the selection process actually makes good progress; but any old physical theory makes idealizations too. And the process under consideration is *interpreted* by us: we assume that it 'selects' for (increase in) the target quality. I don't know whether this sets it apart from ordinary theories. What is decisive, I suppose, is whether we get holism and indeterminacy. Do we?

Biology: the importance of the environment

At least in the case of biological evolution we clearly cannot make any predictions about the fate of a particular population of organisms (a particular set of tokens) without taking into account that population's environment – a holistic aspect. Transfer the population into a different environment,[10] or modify the given one, and selection pressures will change, that is, the web of functional relations between the possible traits and abilities of organisms will shift in many places, such that some qualities that formerly were good for reproductive success (derivatively 'good') are now bad for it, and vice versa. On the other hand, the overarching objective of maximizing reproductive prowess will stay the same, so there is no question of which target quality is the right one to ascribe now – no indeterminacy in that regard. *Unless* what we ascribe isn't just the web's pinnacle but the network as a whole. But that would be as if upon adopting the intentional stance, we ascribed to a given intentional system not just beliefs and desires but the whole network detailing what is a reason for what.

Changing reasons for selection: indeterminate?

However, we do make diagnoses from the selection stance about why traits are selected for, and these diagnoses – ascriptions of rationales – may be candidates for indeterminacy: As the environment changes, the reasons why one and the same trait is selected for can shift. Perhaps good eyesight was once important in P mainly for reasons of predator avoidance, but lately its role in mate selection has become rather more important. Then there would have been a period of transition, and in that period the answer to the question of what was the main reason good eyesight was selected for was perhaps indeterminate. (Or it was simply, "Half for avoiding predators, half for choosing mates.") Or perhaps some indeterminacy lurks in questions concerning single organisms, single tokens, and why they were 'evaluated' the way they were? I don't know.

So I am unsure whether the 'selection stance' is really a proper Dennettian stance. I am certain only that intuitively it feels like a natural intermediate between the physical and the design stance.

Why go intentional?

In search of original intentionality?

What motivates Dennett's claim that in order to understand evolution we must adopt the intentional stance? Reading his "Evolution, Error, and Intentionality" (1987a), one might get the impression that Dennett needs Mother Nature as the originally intentional basis from which all other intentionality – for example, ours – is derived. In that paper he *opposes* 'the doctrine of original [or intrinsic] intentionality', according to which

> no artifact could have the sort of intentionality we have. Any computer program, any robot we might design and build, no matter how strong the illusion we may

[10] Not *too* different, or evolution will come to a stop because the organisms will all die – a drastic and final form of 'selection' not covered by the selection stance, as it leaves no 'best' type behind. Actually of course the environment for individual organisms changes all the time, at least from generation to generation, because as evolution goes on the competition keeps changing: usually, what is relatively 'good' today will be mediocre in a few generations' time.

> create that it has become a genuine agent, could never be a truly autonomous thinker with the same sort of original intentionality we enjoy. (Dennett 1987a, 290)

Our own intentionality is derived

But our own intentionality is of course owed to natural selection: we ourselves (more precisely, our bodies) are "artifacts, in effect, designed over the eons as survival machines for genes that cannot act swiftly and informedly in their own interests" (Dennett 1987a, 298).

> We may call our own intentionality real, but we must recognize that it is derived from the intentionality of natural selection,[11] which is just as real—but just less easily discerned because of the vast difference in time scale and size.
>
> So if there is to be any original intentionality—original just in the sense of being derived from no other, ulterior source—the intentionality of natural selection deserves the honor. What is particularly satisfying about this is that we end the threatened regress of derivation with something of the right metaphysical sort: a *blind* and *unrepresenting* source of our own sightful and insightful powers of representation. (Dennett 1987a, 318)

No need for (strongly) original intentionality

While I agree that the ultimate source of intentionality must be 'blind and unrepresenting', I do not see why it should be intentional in any sense at all. In other words, I do not see any regress of derivation threatening: why must intentionality be either original (in a strong sense that implies, e. g., determinacy of interpretation and first-person authority about meanings; see Dennett 1987a, 300, cf. Goldman 1988) or derived (from someone else's intentionality)? It is much more plausible to assume that it is *our* intentionality which is 'original' – in the weak sense of "being derived from no other, ulterior [itself intentional] source".

Does 'design' presuppose intentionality?

If, however, the design stance were to presuppose the intentional stance, in the sense that the postulation of a designing rational agent is an integral part of the design stance, then a derivation regress for intentionality seems indeed unavoidable: Any rational agent, that is, any intentional system, is itself a prime example of 'design' and so would imply the existence of a further intentional system which is responsible for this design. If Mother Nature is an intentional system (or even just a designed system) then who was this system designed by?

Reality of rationales

So, while it might at first glance seem as if Dennett wants natural selection to be an intentional system because he believes that there must be original intentionality somewhere, it is clear from his writings that he doesn't credit *strongly* original intentionality anyway. Then why doesn't he stick to the weak originality of *our* intentionality and just give up on Mother Nature's? A more plausible motivation for his defense of Mother Nature's intentionality is that he wants evolutionary rationales – 'what Mother Nature had in mind' (Dennett 1987a, 299) – to be real, and that he believes this to require the existence of a rational agent 'appreciating' them: Mother Nature.

> While it can never be stressed enough that natural selection operates with no foresight and no purpose, we should not lose sight of the fact that the process of

[11] In another place, however, he says that "our intentionality is derived from the intentionality of our 'selfish' genes! *They* are the Unmeant Meaners, not us!" (Dennett 1987a, 298; cf. Goldman 1988). The intentionality of our genes is not quite the same as that of natural selection. Both are useful metaphors; both must be taken with many grains of salt; and both are, I think, not necessary for understanding evolution (for genes, the design stance suffices). Importantly, the 'interests' of our genes do not coincide with those of evolution: each gene 'wants' us to be good at propagating *it*, if necessary at the expense of other genes, whereas Mother Nature 'wants' us to be good at propagating *all* our genes. But this needn't make the two intentional stories incompatible; maybe they are just complementary ways of telling one big story about how our intentionality arose.

> natural selection has proven itself to be exquisitely sensitive to rationales, making myriads of discriminating "choices" and "recognizing" and "appreciating" many subtle relationships. To put it even more provocatively, when natural selection selects, it can "choose" a particular design *for one reason rather than another,* without ever consciously—or unconsciously!—"representing" either the choice or the reasons. (Dennett 1987a, 299)

I can't see, however, why in order to be real the rationales behind 'selecting' certain traits or behaviors – the functional interrelations between more or less derivatively and ultimately 'good' qualities – should need agents having 'appreciated' them and 'acted' upon them. To be a real pattern, it seems to me, all that is needed is that the intended (type of) phenomenon could in principle be apprehended, and this appreciation be put to use, by a suitable possible type of (finite, physical) agent. With respect to the patterns instantiated by natural selection we are such agents, and even if no agents like us had ever evolved, our mere possibility already suffices.

Closing remarks

Strength of my arguments?

I have argued that the intentional and the design stance are inappropriate for understanding evolution, but I am afraid that my arguments to that effect are at best persuasive, not compelling. However, I hope that the selection stance I have proposed introducing is, though unfamiliar, more plausible as an explanatory strategy for evolution than the two other stances, and that this constructive suggestion will add weight to my destructive thesis. To me at least it seems extremely plausible that it is this stance which comprises the right amount, and kind, of structure for understanding natural selection.

The selection stance: minimal normativity?

A bare suggestion: Maybe the selection stance captures a very abstract and general, perhaps minimal, form of *normativity*. It considers a variety of options (types to be tokened), which are 'encouraged' or 'discouraged' (selected or rejected), according to some conception of goodness, by their instances engendering 'reward' or 'punishment' (positive or negative evaluation). Does normativity make sense without some force ('selection pressure', in the case of evolution) operating to increase the 'good'? If the selection stance is indeed applicable, that is, if the selection process really works, after some time the better options will be realized more frequently and the worse options less frequently. It is questionable, however, how well the selection stance is applicable to real-life instances of normativity (besides design/adaptation and rationality), for example, to moral or aesthetic normativity, where improvements over time may be rather meager.

References

Dawkins, Richard. 1976. *The Selfish Gene*. New York/Oxford: Oxford University Press. Second edition, with two additional chapters, 1989.

———. 1996. *Climbing Mount Improbable*. New York: Norton.

Dennett, Daniel C. 1971. "Intentional Systems." *Journal of Philosophy* 68 (4): 87–106. Reprinted in his *Brainstorms: Philosophical Essays on Mind and Psychology*, Cambridge Mass./London: MIT Press (A Bradford Book) 1978, pages 3–22.

———. 1981a. "Three Kinds of Intentional Psychology." In *Reduction, Time and Reality: Studies in the Philosophy of the Natural Sciences*, edited by R. Healy, Cambridge:

Cambridge University Press, pages 37–61. Reprinted in Dennett 1987b, 43–68; references are to the reprint.

———. 1981b. "True Believers: The Intentional Strategy and Why It Works." In *Scientific Explanation: Papers Based on Herbert Spencer Lectures Given in the University of Oxford*, edited by A. F. Heath, Oxford: Clarendon Press, pages 53–75. Reprinted in Dennett 1987b, 13–35; references are to the reprint.

———. 1983. "Intentional Systems in Cognitive Ethology: The 'Panglossian Paradigm' Defended." *Behavioral and Brain Sciences* 6 (3): 343–55. Reprinted in Dennett 1987b, 237–68; references are to the reprint.

———. 1987a. "Evolution, Error, and Intentionality." In Dennett 1987b, 287–321.

———. 1987b. *The Intentional Stance*. Cambridge Mass./London: MIT Press (A Bradford Book).

———. 1988. "Science, Philosophy, and Interpretation." *Behavioral and Brain Sciences* 11 (3): 535–44.

———. 1991. "Real Patterns." *Journal of Philosophy* 88:27–51. Reprinted in his *Brainchildren: Essays on Designing Minds*, Cambridge Mass.: MIT Press (A Bradford Book) 1998, pages 95–120.

———. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon & Schuster.

Dretske, Fred. 1988. "The Stance Stance." *Behavioral and Brain Sciences* 11 (3): 511–12.

Goldman, Alvin I. 1988. "Derived Intentionality?" *Behavioral and Brain Sciences* 11 (3): 514.

von Bülow, Christopher. 2002. "Cui bono? Evolution, Selektion und Interessen." Manuscript at http://www.uni.kn/FuF/Philo/Philosophie/philosophie/files/interest.pdf.

———. 2013. "Mem." In *Enzyklopädie Philosophie und Wissenschaftstheorie*, second edition, Volume 5, edited by J. Mittelstraß, Stuttgart/Weimar: J. B. Metzler, pages 318–24. English translation at http://www.uni.kn/FuF/Philo/Philosophie/philosophie/files/meme.pdf; references are to the English version.