

Penultimate draft

# Behavioral Traits, the Intentional Stance, and Biological Functions: What Neuroscience Explains

Marcel Weber  
Department of Philosophy  
University of Konstanz  
marcel.weber@uni-konstanz.de

## **Abstract**

It has been claimed that the intentional stance is necessary to individuate behavioral traits. This thesis, while clearly false, points to two interesting sets of problems concerning biological explanations of behavior: The first is a general in the philosophy of science: the theory-ladenness of observation. The second problem concerns the principles of trait individuation, which is a general problem in philosophy of biology. After discussing some alternatives, I show that one way of individuating the behavioral traits of an organism is by a special use of the concept of biological function, as understood in an enriched causal role (not selected effect) sense. On this view, a behavioral trait is essentially a special kind of regularity, namely a regularity that is produced by some regulatory mechanism. Regulatory mechanisms always require goal states, which can only be provided by functional considerations. As an example from actual science, I examine the case of social behavior in nematodes. I show that the attempt to explain this phenomenon neuroscientifically actually transformed it. This supports the view that scientific explanation does not explain an explanandum phenomenon that is given prior to the explanation; rather, the explanandum is changed by the explanation. This means that there could be a plurality

of stances that have some heuristic value initially, but which will be abandoned in favor of a functional characterization eventually.

### **1. Introduction: The Intentional Stance and the Observation of Animal Behavior**

According to some philosophers, we take an "intentional stance" towards animals, especially but not exclusively towards rational ones (Dennett 1989). This means that we bring some of their activities and movements under intentional concepts, i.e., concepts that involve an attribution of intentional states to the entity whose behavior we want to explain. Intentional states present some object as *being so*, for example, an incoming object as being a dangerous predator. Even though we may not have any direct evidence for the presence of such states, treating living beings *as if* they had such states allows us to predict them much more easily. Consider a gazelle that is being hunted by a leopard. Attributing to the gazelle intentional states to the effect that it *thinks* that the charging object is a dangerous cat, or at least something as dangerous as a snake but much, much faster, makes its ensuing behavior more predictable. The point is not that attributing intentional states to the gazelle is the best *explanation* of its behavior, it is merely that taking some other stance toward it, for instance, the stance that a physicist takes towards two moving masses, make its movements much more difficult to predict for us. Let's refer to this as the *predictive value of the intentional stance*.

Here is a second, related idea that is sufficient enough to merit separate consideration. Looking at animals with our intentional goggles on unveils certain *general patterns* that would otherwise remain invisible. This thesis has been defended by Hilary Kornblith in his *Knowledge and its Place in Nature* (Kornblith 2002). Kornblith considers reports from ornithologists such as this one:

"George Schaller told me of watching raven pairs in Mongolia *cooperate* in snatching rats from feeding raptors. Similarly, in Yellowstone Park, Ray Paunovitch reported seeing a re-tailed hawk with a ground squirrel. Two ravens approached. One *distracted* the hawk from the front while the other handily snatched the squirrel from behind. Carsten Hinnerichs saw the same maneuver repeated three times in a row in a field near Brücke, Germany, where a fox was catching field mice. Terry McEneaney, Yellowstone Park ornithologist, observed two ravens circling an osprey nest where the female osprey was incubating. One raven landed on the nest rim and took a fish, then while the osprey was *distracted* by this thief, the other ravens swooped down and stole an osprey egg." [p. 31 emphasis mine]

Kornblith emphasizes in particular that the use of intentional notions such as "cooperate" and "distract" allows us to recognize what quite different behaviors in anatomically and behaviorally quite different species have in *common*. For example, he draws attention to the fact that in the description of the "distracting"-behavior, the ornithologists do not even mention any bodily motions:

"Heinrich does not speak of the ravens moving their beaks, or wings or bodies in certain patterns. Indeed, there is no reason to think that the manner in which the ravens distracted their various targets involved any commonality at all at the level of bodily motion. There are bits of animal motion that may be described in such terms, but this does not seem to be such a case. Instead, what is common to the various episodes described can only be appreciated by attributing certain intentional states to the animals involved. If we see the behavior as a case of one bird distracting

another, we are able to make sense of it in a way that a description in terms of moving beaks, wings, and bodies fail to capture.” [p. 33]

The point is not that without the intentional stance, there are no patterns discernible at all. Kornblith allows descriptions such as “wing-flapping, squawking, pecking, and so on” as descriptions that are entirely free of intentional notions. The point is rather that, if viewed in non-intentional terms, these patterns are *heterogeneous*. Only an intentional category such as “distraction” allows us to view these patterns as instances of a single kind, a kind that also includes behaviors in animals that lack wings and beaks, and are utterly bad at squawking.

Thus, Kornblith's contention is not that animal behavior can only be *explained* in intentional terms. The point is rather that it is more often than not *described* in such a way, before anyone even ventures to propose an explanation.

Kornblith goes as far as to suggest that his point stands even for behaviors that seem to admit of fairly simple descriptions, because they can be described without any reference to complicated movements of the part of an animal relative to each other, but simply with reference to the spatio-temporal location of the entire animal. The examples he discusses here are mostly animals that are capable of homing.

Thus, Kornblith's point is that the intentional notions allow us to bring patterns that are quite heterogeneous at the physical level under *general kinds*. Let us refer to this property as the *classificatory value* of the intentional stance.

The goal of Kornblith's argument is to show that intentionality and cognate notions pick out perfectly fine natural kinds that are on a par with better-known natural kinds from science, such as gold or oxidation-reduction reactions.

Kornblith's use of the ethological literature is quite objectionable; because he does not check what role the reports he cites actually play in these investigations. It's often a mistake to take scientific language too literally. Nonetheless, this kind of arguing does raise some interesting issues in philosophy of science, namely the problem of the theory-ladenness of observation and the question of what constitutes a good explanandum for (neuro-) biological explanations of behavior in the first place.

In this chapter, I want to address these issues in the context of a recent example from neuroscience, namely social behavior in the nematode *C. elegans*. In Section 2, I shall motivate the problem of trait individuation a little further and critically examine some arguments from the literature. This will make clear that trait individuation in biology is often based on the concept of function. In Section 3, I discuss a concept of function that I claim is suited for this task. In Section 4, I present my central neurobiological example and show that it illustrates functional trait individuation. Furthermore, I show that research on this model organism actually transformed what the neuroscientists

## **2. The General Problem of Trait Individuation**

Kornblith's point is an instance of what philosophers of science refer to as the theory-dependence or more metaphorically the "theory-ladenness" of observation. The point is that ornithologists' reports of animal behavior do not merely deploy pure observation terms – even if there should be

such terms. Rather, these reports employ theoretical notions. The theory that supplies these notions, according to Kornblith, is *folk psychology* (this is my term, not Kornblith's). For notions such as "cooperation" and "distracting" are precisely those that we use in our everyday interactions with our fellow human beings, and sometimes also with animals. (It is a good idea to try to "distract" a bear that is approaching a small child. It is a very, very bad idea to try to push the bear away by brute force).

Now it is possible to accept the claim that descriptions of behavior in biology are theory-dependent while, at the same time, rejecting the idea that the theory that such descriptions depend on is folk psychology.<sup>1</sup> Of course, such a reply would have to explain why intentional vocabulary (while it was banned from both psychology and ethology for a long time, namely from the heydays of behaviorism until quite recently) is so widely used in the scientific literature, and why cognitive ethology is flourishing to the extent that it has been in recent years. However, it should be noted that cognitive ethology is precisely the field that problematizes accounts of animal behavior that involve intentionality, so it may not be representative.

---

<sup>1</sup> That certain ways of classifying animal behavior are theory-laden has also been pointed out by Philip Kitcher in *Vaulting Ambition*, his powerful critique of sociobiology (Kitcher 1985). He argues that the way in which certain behavioral traits such as rape are singled out for evolutionary explanations is often informed by ideological preconceptions and false stereotypes, e.g., that rape is sexual behavior.

Before I take on the task of developing an alternative view as to what theoretical notions (ought to) inform behavioral observations, I shall briefly lay out the general problem of trait individuation as it has been discussed in the philosophy of biology.

The problem we are dealing with is an instance of the more general problem of how the traits of organism are individuated. By "traits" I do not mean only genetic traits, not even traits that vary in a population. The question is rather how certain parts of an organism such as organs are identified as some sort of a unified structure in the first place. While this may seem obvious in some cases like the leaves on the beech tree in front of my window, it is difficult to state some general or even medium-range principles as to what is a good way of cutting up an organism along its natural joints. Clearly, Socrates's advice to proceed by trying not to splinter any parts, "as a bad butcher might do" (Plato 1997, *Phaedrus*, 265e) is not helpful at all, for we have no theory-independent way of knowing when we have splintered something. One way of putting our problem is thus whether he or she who cuts behaviors with intentional knives is a good butcher or a bad one. The case is particularly difficult in behavioral biology, as behaviors can involve different parts of an organism, and may consist of non-contiguous sequences of events that spread out in time.

What are the options? What are the conceptual tools used by a good butcher of animal behavior? While there is hardly any literature on this specific question, some philosophers of biology have addressed the problem of trait individuation in general. They have come up with two kinds of answers: *Proper functions* and *homology*.

An adherent of the first answer is Alex Rosenberg (Rosenberg 2006). He argues that parts of organisms are individuated via their biological functions, which he understands in terms of proper functions. On this view, some item X has a function F in organism S exactly if X does F and the fact that some earlier tokens of X have done X is a cause of X's presence in S. The way in which earlier tokens can cause the presence of some item in later generations, of course, is natural selection. Thus, Rosenberg's view is that natural selection is not only needed to explain why some organism S came to have a part X, but to speak of X as some kind of unity in the first place. It is for this reason that Rosenberg thinks that the theory of natural selection is fundamental for the whole of biology. This, of course, includes behavioral biology. According to Rosenberg, the description of behavioral traits is laden and/or ought to be laden by theoretical hypotheses about selection history. A trait such as a wing is individuated by the fact that it was selected for flying, no matter what other capacities it may have (for instance, it's capacity of being flapped so as to distract some other animal). On this view, descriptions of an organism's trait are laden by the theory of natural selection and assumptions about the evolutionary past.

Paul Griffiths (forthcoming) has argued that this view puts the cart before the horse. The parts of organisms and their causal capacities must be understandable independently of natural selection.

Otherwise, the following regress threatens:

1. Selected effect functions are ascribed by causal analysis of the capacities of the parts of ancestral organisms and a determination of their fitness contribution.
2. Thus, we must already be able to individuate the parts. This cannot be done on the basis of the ancestors to the ancestral organisms, because it would generate a regress



3. But if we are able to individuate parts for ancestral organisms independently of their selection history, then this is possible for living organisms

[This is modified from Griffiths, forthcoming]

So if natural selection is not fit for the individuation of organismic parts, what is? One possibility is another fundamental concept of biology, that of *homology* (Griffiths 2006)(Love 2007)(Brigandt 2007). Richard Owen defined homologues as "the same organ in different animals under every variety of form and function". For example, a bird's wings are said to be homologous to our arms, while the panda's famous "thumbs" are not homologous to our thumbs. Obviously, Owen's definition is purely formal; it does give us any criteria when two organs count as "the same". It merely says that homology is a sameness-of-kind relation that abstracts away from form *and* function. The standard view of homology sees it as a relation of *shared ancestry*. However, this view has some known difficulties.

The main difficulty consists in saying what it means for a *part* of an organism to have the same ancestry as some part of another organism. Parts of organisms do not reproduce, at least not independently of the individual of which they are a part. Thus, it seems that, not unlike in the case of selected effect functions, biologists must *first* apply some *other* concept to individuate characters and *then* trace that part through the phylogenetic tree via the ancestry of the organisms that have these parts. Some have suggested that this elusive other individuating property might be something like developmental units, however, this has turned out to be difficult to spell out. At any rate, this seems highly unsuitable for behavioral traits, which are normally not at all viewed in developmental terms.

The issue of homology is, of course, extremely complex. My considerations here are only meant to show that it is unlikely that behavioral traits normally often individuated by homology.

So now we have considered the following candidates of theories or theoretical notions that have been claimed to inform – for better or for worse – the classification of behavioral traits: (1) folk psychology, (2) ideology, (3) natural selection, (4) the theory of descent and modification, (5) theories about development. It is not my goal here to provide a critical assessment of these various possibilities. Rather, what I would like to do is to sketch an alternative.

This alternative consists of two movements. The first movement will be to grant biological functions an important role in the individuation of behavioral traits. However, the salient concept of function will not be that of proper functions, but a modified version of the concept of causal role functions (Section 3). The second movement will consist showing that the investigation of a behavioral trait can actually transform the explanandum. This will be shown on an actual example (Section 4).

### **3. Biological Functions**

I have argued that a construal of functions in terms of selected effect functions is unfit for the task of trait individuation. As an alternative, I suggest a modified version of causal role functions ((Weber 2005a) (Weber 2005b)). This account starts with Cummins's (1975) analysis according to which functions are such capacities that are capable of explaining a capacity of some containing system. The paradigm is the heart's capacity to pump blood figuring in any adequate explanation of the circulatory system's capacity to transport nutrients, oxygen and blood cells

through the body. According to Cummins, the pertinent capacity of the containing system is a matter of an interest-based choice to be made by the investigator. I have modified this account by suggesting that this systems capacity should be made dependent not on the investigator's interests, but on the role that the containing system itself plays in the self-reproduction of the *whole* organism (see (Davies 2003) for a similar idea). I argue that this is what turns Cummins-functions into *biological* functions. Cummins-functions can be applied to any kind of system. But only biological systems are capable of self-reproduction. In order for self-reproduction to occur, an organism's functions must work together. The specific contribution that some organ's causal capacities make to self-reproduction makes will depend on what other organs do. For example, if there were subsystems of an organism that would use the heart's heat production towards something that itself makes a contribution of self-reproduction, then the heart would (also) have the function of producing heat. It is the place that such a causal capacity plays in a whole network that gives it its function (perhaps much in the way in which a linguistic expression's meaning is given by the inferential role that the expression plays in a network of other expressions, as claimed by inferentialists and semantic holists).

I have argued that introducing such a global constraint on a system of functions might make the interest-dependence vanish, provided that there is exactly one way of laying a network of cooperating functions over an organism. Of course, this is hard to prove; but I suggest that it might be possible by using a notion of maximal explanatory coherence (Weber 2005b).

I have argued in my (2005a) that it is such a concept of function that underlies much of experimental biology. I would now like to suggest that such a concept also underlies the

individuation of behavioral traits, at least in some cases. To demonstrate this, I shall finally turn to a real example.

#### **4. Example: Social Behavior in *C. elegans***

The example comes from research on the soil nematode *Caenorhabditis elegans*. This is a tiny worm, about 1mm long, which feeds on bacteria. Even though its nervous system is extremely simple, consisting of only 302 neurons, this roundworm exhibits a form of social behavior: If placed on a bacterial lawn (which is their favorite food), the worms clump together for feeding. The functional significance of this is not entirely clear until this day. It seems that this behavior is somehow regulated, as it does not occur when the bacteria are spread (as opposed to forming a lawn) or when there are no bacteria at all. Furthermore, the worms do not merely move about randomly and then get stuck to each other. It seems that they move towards each other and respond to the vicinity of clumps by slowing down. What is striking is that there exist strains of the worm that do not clump together at all. Worm scientists Mario de Bono and Cornelia Bargmann showed that this difference can be causally attributed to a single gene, dubbed *npr-1*, for which there is natural variation (de Bono und C. I. Bargmann 1998). The gene encodes a receptor for a small peptide that showed a high sequence similarity to a protein also found in other animals including humans: the so-called neuropeptide Y. This peptide is a neurotransmitter, in other words, it is involved in the transmission of signals between neurons. Most interestingly, it seems to be involved in the regulation of food intake in mammals, including humans.

What lessons can we draw from this example? The most important one, in my view, is this:

*The classification of this phenomenon as behavior already involves bringing it under functional concepts. If the worms would just move around randomly and get stuck to each other, it would not be classified as behavior but as some sort of surface adhesion phenomenon that interferes with the worms' mobility. To speak of behavioral phenomenon implies the existence of a specific regulatory mechanism that responds to certain environmental stimuli (here: bacterial lawns and conspecific worms) and produces a functional response.<sup>2</sup>*

By bringing the clumping of worms under the concept of behavior, the biologists are not yet committed to what the actual function of this regulatory mechanism is. In order to attribute a specific function, more knowledge is required, at least according to the account of functions that I have briefly summarized in the last section.<sup>3</sup> All that is involved in seeing a behavioral trait here is that there exists some hypothetical regulatory mechanism that has *some* biological function.

So far, it is not clear that we are already dealing with a well-individuated trait. But at the very least, it is distinguished from other kinds of phenomena that might look similar where animals clump together for purely accidental reasons. But it might not yet be conceptually distinguished from other forms of animal clumping.

---

<sup>2</sup> Some philosophers such as (Millikan 1984) would already see the operation of such a control mechanism as a simple form of intentionality (provided that it is an evolutionary adaptation), but that's controversial, to say the least.

<sup>3</sup> It is not necessary to understand all of an organism's functions on my systemic account. It is enough to know some of the next nodes in a function graph. Here, in all likelihood the most salient systems capacity will be foraging for food.

Another question is what warrants the classification of this behavior as 'social'. It should be quite obvious that it is social in quite a weak sense, namely in the sense that the worms regulate their feeding behavior such as to functionally respond to the presence of conspecifics. More elaborate forms of social behavior will involve recognition and differential responses to different individuals (which enables certain solutions to game theoretic situations like the infamous 'tit-for-tat' solution to the prisoner's dilemma). Even more elaborate forms will involve a theory of mind, self-consciousness and, as some people think, collective intentionality. But my point here is that the classification of certain phenomena as *behavior* in the first place involves the concept of biological functions rather than an intentional stance or any of the other options that we have discussed. Thus, the description of behavioral phenomena is sometimes laden with functional notions.

I do not want to make any strong claims to the effect that this is always the case, or that this is the One Right Way of individuating behavioral traits. My reasons for this are not so much a general pluralistic stance, but another fundamental insight into the classification and explanation of phenomena in science. This insight is found, along with the general idea of theory-ladenness, in Paul Feyerabend's classic article "Explanation, Reduction, and Empiricism" (Feyerabend 1962). It is time to put a spotlight on the salience of this insight for the practice of biology, especially behavioral biology.

The insight is that scientific explanations of a phenomenon do usually not leave the explanandum phenomenon the same. In other words, the following picture will be rejected: That biologists can first pick out some pattern or phenomenon for explanation and then provide an explanation for

that phenomenon, for example, some mechanism or some adaptive evolutionary story. Rather, the phenomenon to be explained is *transformed* in the very attempt of explaining it.

In our example from worm biology, the inquiry into the underlying mechanisms of the "social feeding" behavior in *C. elegans* made the biologists realize that they only had the elephant's trunk (or perhaps its tail, or ...) when they started to investigate the phenomenon. The following additional findings from the worm study mentioned should make this clear.

De Bono and Bargmann used the isolated *npr-1* gene to determine in which strains it was present and absent, respectively. They compared 17 different strains, 15 of which were natural isolates from different parts of the world. These strains fell into two groups: those that exhibit social feeding and those that do not. Remarkably, this variation could be fully accounted for by a single amino acid substitution in the *npr-1* gene. The distribution of these alleles provided reasons for thinking that this mutation arose only once in evolution. No matter which allele was ancestral, both alleles were not only able to be maintained, but to actually spread to different parts of the world and quite different habitats. In some locations, both kinds of strains (social and solitary) were found coexisting. This is strong evidence for the hypothesis that both variants are *adaptive* under at least some conditions. Thus, the so-called "solitary" are not somehow defective; they merely have a different reaction norm to some environmental cues.

It should also be mentioned that when bacteria do not form a lawn but are so scarce as to become growth-limiting, all strains exhibit the clumping, and in the total absence of bacteria all of them move around solitarily. There are also indications that different worms respond differently to different bacterial species, which they may recognize by the chemical profiles of their

metabolites. De Bono and Bargmann conclude from these findings: "It is likely that the behaviors we observed represent a narrow view of a more complex behavioral choice."

The "complex behavioral choice" alluded to by the worm biologists may be an intricate regulatory mechanism that chooses between different behavioral repertoires depending on the availability of and quality of food and the density of worms. An important piece of evidence for this hypothesis is the fact that the neuropeptide receptor encoded by the *npr-1* gene seems to be involved in the regulation of food intake and metabolism in a wide variety of animal species from extremely remote taxonomic groups. It is also thought that, in the worms, it activates the formation of so-called *dauer*-forms of the worm, which is a dormant state that allows the worms to survive starvation. There also seem to be connections to the pheromone response.

The case thus allows us to draw another lesson:

*As a result of the investigation, the behavioral phenomenon that the scientists want to explain changed: It started as a functional response to bacterial lawns and other worms, namely clumping. Now we are dealing with a more complex phenomenon that includes not just clumping but also swarming, the onset of reproduction and the formation of the dormant state. What seemed to be different phenomena at first are now treated as one phenomenon. What unified them was a combination of causal knowledge, extrapolating gene-function relations from other organisms, knowledge about biogeography and the autecology of the worms, and the pursuit of other questions raised by the initial findings.*



None of the initial findings really reveals anything about the underlying mechanism of the worms' behavioral response to the availability of food and the presence of conspecifics. Perhaps even more remarkable is the fact that it is not yet entirely clear what the functions of the initially observed clumping behavior actually are. One family of hypotheses says that the function of clumping must have something to do with the fact that it creates a microenvironment with a reduced oxygen level, perhaps to avoid oxidative stress (Gray u. a. 2004). Alternatively, it might be a side-effect of the animals seeking a low-oxygen environment because that's where the food normally is (anaerobic bacteria). That way, they might also end up in clumps. So it is not clear that the clumping itself has a function.

This is not in contradiction to my earlier claim that to classify the clumping phenomenon as behavior involves an application of the concept of function. The hypothesis that generates this classification is weaker, perhaps something like this: There exists a regulatory mechanism  $M$  with respect to a function  $F$  that explains the occurrence of the worm clumps under conditions  $C$ . Note that regulatory mechanisms are necessarily connected to some function, otherwise they would not be *regulatory* mechanisms. Regulation always needs a set of *goal states*, which can only be states that have some functional significance for the organism (on pains of applying the notion of regulation in its original, intentional sense; see (Canguilhem 1988)).

## **5. Conclusions: A Plurality of Stances?**

If the modification of behavioral phenomena by the very attempt of explaining them occurs regularly, we might be able to give a more pluralistic answer to our question of what is a good way of individuating behavioral traits or, more generally, for carving out behavioral phenomena. There may not be One Right Way of doing this. If the initial classification will change anyway as

the underlying causes are unveiled, perhaps it is perfectly all right to start the classificatory task with folk psychology, or with some other stance. But my suspicion is that, as they learn more about a phenomenon that was initially picked out under the intentional stance, biologists will quickly move towards using functional concepts for classifying behaviors, as I demonstrated it for the worm case. Folk psychology may have a great predictive value, when it comes to naturally classifying the characteristics of organisms it seems wiser to use functions. For what could be more biological than that?<sup>4</sup>

Such an account of trait individuation is needed to complement existing accounts of explanation in neuroscience, such as Carl Craver's (2007) ground-breaking account. Craver gives an excellent analysis of what it means to explain some (neuro-) biological phenomenon. A phenomenon consists of the regular behavior of some biological entity, and explaining it consists in exhibiting a mechanism that produces this behavior. Craver brilliantly analyzes the conditions such a mechanism must fulfill. However, he doesn't say much about what constitutes a phenomenon that is worth explaining in the first place. He does show that mechanisms sometimes form inter-level hierarchies (see especially Chapter 5). These hierarchies often top-off at the level of behavior and bottom-out at the level of atoms. But what makes biologists pick some top-level behavior as their explanandum in the first place? To this question, I have tried to give an answer here. It should be noted that the significance that functionally individuated behavioral traits bestow on the

---

<sup>4</sup> I don't mean to suggest that this is a general solution to the problem of trait individuation. As Ken Waters has convinced me, there is no general solution to this problem. Trait individuation depends on context. For example, sometimes biologists will pick traits precisely because they have no function and are thus neutral with respect to fitness and natural selection. I wish to thank David Sloan Wilson for bringing this kind of example to my attention.

mechanisms that explain them trickles down the inter-level hierarchies that are characteristic of mechanistic explanation in neuroscience and elsewhere in biology.

## Acknowledgements

I wish to thank Daniel Sirtes, Katie Plaisance and Thomas Reydon for helpful comments.

## References

- de Bono, M., and C. I. Bargmann. 1998. "Natural variation in a neuropeptide Y receptor homolog modifies social behavior and food response in *C. elegans*." *Cell* 94(5):679-89.
- Brigandt, Ingo. 2007. "Typology now: homology and developmental constraints explain evolvability." *Biology and Philosophy* 22(5):709-725.
- Canguilhem, Georges. 1988. *Ideology and Rationality in the History of the Life Sciences*. Cambridge, Mass.: MIT Press.
- Craver, C. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- Davies, Paul Sheldon. 2003. *Norms of Nature: Naturalism and the Nature of Functions*. The MIT Press.
- Dennett, Daniel C. 1989. *The Intentional Stance*. The MIT Press.
- Feyerabend, Paul K. 1962. "Explanation, Reduction and Empiricism." In *Scientific Explanation, Space, and Time. Minnesota Studies in the Philosophy of Science, Vol. III*, Vol. III: University of Minnesota Press, p. 28-97.
- Gray, Jesse M. u. a. 2004. "Oxygen sensation and social feeding mediated by a *C. elegans* guanylate cyclase homologue." *Nature* 430(6997):317-322.
- Griffiths, Paul. "In What Sense Does "Nothing Make Sense Except in the Light of Evolution"?" *Acta Biotheoretica*.
- Griffiths, Paul E. 2006. "Function, Homology, and Character Individuation." *Philosophy of Science* 73(1):1-25.
- Kitcher, Philip. 1985. *Vaulting Ambition: Sociobiology and the Quest for Human Nature*. Cambridge MA: MIT Press.

- Kornblith, Hilary. 2002. *Knowledge and its Place in Nature*. Oxford: Clarendon Press.
- Love, Alan. 2007. "Functional homology and homology of function: biological concepts and philosophical consequences." *Biology and Philosophy* 22(5):691-708.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, Mass.: MIT Press.
- Rosenberg, Alex. 2006. *Darwinian Reductionism*. Chicago: The University of Chicago Press.
- Weber, Marcel. 2005a. "Holism, Coherence, and the Dispositional Concept of Functions." *Annals in the History and Philosophy of Biology* 10:189-201.
- Weber, Marcel. 2005b. *Philosophy of Experimental Biology*. Cambridge: Cambridge University Press.