

ADAPTATION AND MORAL REALISM

William Harms
Centre for Applied Ethics
University of British Columbia
Rev: January 2, 1999
419-372-7219
bharms@interchange.ubc.ca

ADAPTATION AND MORAL REALISM

ABSTRACT: Conventional wisdom has it that evolution makes a sham of morality, even if morality is an adaptation. I disagree. I argue that our best current adaptationist theory of meaning offers objective truth conditions for signaling systems of all sorts. The objectivity is, however, relative to species — specifically to the adaptive history of the signalling system in question. So while evolution may not provide the kind of species independent objective standards that (e.g.) Kantians desire, this should be enough for the practical work of justifying our confidence in the objectivity of moral standards. If you believe morality is an adaptation, you should be a moral realist.

Keywords: evolutionary ethics, moral realism, teleosemantics, naturalism.

Suppose, as many naturalistically inclined philosophers do, that morality is an adaptation — a product of evolutionary process which has been preferred and refined by selection due to its having performed some useful function. What does this bode for morality? The current consensus seems to be that whatever this gives us, it is not moral realism. Most everyone, nonetheless, seems to want to be able to talk like a moral realist, that is, like someone who believes that statements concerning the moral *wrongness* of classes of acts and the moral *justification* of individual acts are made true or false by some set of externally determined standards that bear a unique constitutive relationship to morality. What results from this tension differs from case to case. Philip Kitcher (1985) in summing up E.O. Wilson’s attempt to “biologize ethics”, writes,

To the extent that people can be viewed as maximizing their own inclusive fitness through cooperation with others, apparent conflicts of interest may be diagnosed as situations in which all the parties maximize their inclusive fitness by coordinating their behavior. Yet there are innumerable situations — among them some of the most troubling — in which the reproductive interests of individuals do clash. For these situations pop sociobiology has nothing to offer. There is no higher standpoint than the dictates of the hypothalamus. There is no impartial perspective. (1985, 433f.)¹

¹ See also Kitcher 1993 for more on what’s wrong with E.O. Wilson.

So, according to Kitcher, Wilson's sociobiology fails to give us the kind of standards genuine morality seems to require. George C. Williams, the grand old man of gene selectionism, has an even more pessimistic view of the prospects of naturalizing morality. Famously, he characterizes nature in the form of evolution as "the tenacious and powerful enemy" of all morality and goodness. (1994, 212) He thinks that,

There is no encouragement for any belief that an organism can be designed for any purpose other than the most effective pursuit of this self-interest. (1994, 197)

How then is the naturalist to reconcile the apparent incompatibility of evolution and moral objectivity? The most fashionable response is to insist that culture somehow emancipates us from the dictates of our genes. This is Williams' view, for instance.

The evolutionary process is immensely powerful and oppressive, but, unlike Job's God, it is abysmally stupid. It can reliably maximize current selfishness at the level of the gene, but it is blind to future macroscopic consequences of current action. It does not have the sense to realize that mechanisms evolved for practicing unfair nepotism or making self-seeking deals with other can be subverted in the interests of broad altruism. Thus can Huxley's "ethical nature" emerge incidentally from biological nature. It is one of many examples of evolutionary changes with important future consequences entirely unrelated to the selfishness responsible for the changes. (1994, 209)

The helping hand of the good Samaritan and the motivation for its use raise no question of the malice or power of natural selection. They merely show that this persistent and powerful enemy is a mindless fool. (1994, 210)

Richard Dawkins goes further than most in claiming that we are free of the dictates of culture as well,

We are built as gene machines and cultured as meme machines, but we have the power to turn against our creators. We, alone on earth, can rebel against the tyranny of the selfish replicators. (1989, 201)

I won't belabor the point here, but one finds similar (though usually less colorful) comments in the writings of Elliott Sober, Richard Alexander, Stephen Jay Gould, Daniel Dennett, and so on.

The problem with this way of making morality and evolution compatible is that it is hard to see where the objectivity of standards is supposed to come from. The substantial cultural

relativism that results from insisting that culture is the source of goodness provides the occasion for numerous subtle and/or counterintuitive maneuvers. I will be most directly concerned with bullet-biting line that Ruse & Wilson (1986) and Dawkins (1989/1976) take to the effect that moral objectivity is an illusion foisted upon us by our genes, but my general point has implications across the board. According to this particular line, there is no objective, species independent set of moral standards, but the illusion that there are such standards is generally a good thing. (Good for fitness, that is). Apparently, evolutionary theory shows moral objectivity to be an illusion, our beliefs to the contrary the stratagems of deceitful genes.

I beg to differ. My argument is simple. Apart from pointing out certain confusions regarding the use of "objective" and "illusion", all I want to say is that the current best theory of meaning built on the analysis of language as an adaptation works equally well for bee dances, beaver tail splashes, monkey warning cries, indicative/descriptive human language, and, of course, the systematic and (by presumption) adaptive noises we make as an essential part of our moral life. I refer primarily to Millikan's (1984, 1993) "biosemantics" or "teleosemantics", but also to the work of Dretske and others in which normative biological functions play a role in naturalistic semantics. Truth, on Millikan's version, does not *explain* the value of the true sentence or signal. Rather, truth of a signal just is the fact that the signal now stands in the kind of relation to the world which explains the historical contribution to fitness of the signaling system. There are true and false bee dances, true and false monkey warning cries, and if there is the kind of adaptive history we began by supposing, then there are (objectively) true and false moral statements/imperatives. This truth and falsity does not require "special" moral facts (moral entities and/or moral properties) to serve as the truth conditions for moral statements, nor is this truth reducible to the truth of scientific descriptions of the adaptive history of morality. Happily, however, truth conditions for moral imperatives and claims of justification can in principle be

specified in scientific terms. Of course, whether any of this is true is an empirical question, with much of the evidence lost in the mists of history. My point is that, if you believe that morality is an adaptation, then you ought to believe that there objective standards that apply, and that there is a distinctively *moral* kind of truth.

Objectivity, Species-Independence, and Illusions

One common misapprehension is that objective standards must be species independent.

... human beings function better if they are deceived by their genes into thinking that there is a disinterested objective morality binding upon them, which all should obey. ... In short, ethical premises are the peculiar products of genetic history. And they can be understood solely as mechanisms that are adaptive for the species that possess them. It follows that the ethical code of one species cannot be translated into that of another. No abstract moral principles exist outside the particular nature of the individual species.

It is thus entirely correct to say that ethical laws can be changed, at the deepest level, by genetic evolution. This is obviously quite inconsistent with the notion of morality as a set of objective, eternal verities. (Ruse & Wilson, 1986)

Clearly, whatever functional norms attach to biological adaptations, they are not species independent, but rather contingent on the particular adaptive history of the faculty in question. From this it is supposed to follow that they are not objective standards. This view is, of course, consistent with the usual approach to moral objectivity. Kant was quite insistent that his imperatives were not derived from anthropology, but from the nature of reason and the will. Rational Martians were just as bound as we are. Peter Singer's (1975) extension of utilitarianism to animal rights makes clear that being human is not a prerequisite to consideration in the Benthamite calculus, it is pleasure and pain that are at issue. Any time the principles that attempt to ground objective moral norms do not make reference to human nature, the standards cannot but be species independent.

Be this as it may, this does not mean that standards that apply to every human being everywhere simply because they are *human* beings are not objective in a sense adequate for moral

realism. Human beings are equipped with a variety of adaptations, some shared with other species, some unique. To have an adaptation is to be capable of malfunction. Such malfunction consists not in failing to maximize inclusive fitness now, but in failing to perform in the ways which the feature in question has historically contributed. The standards for the proper functioning of an adaptation are objective, since they derive from the actual history of the feature in question. If all human beings, due to common ancestry, share a particular adaptation, then there is a historically rooted set of standards that apply to them all. If the adapted systems in question are flexible enough, then it may be that different environments/cultures make different configurations or behaviors "proper". Nonetheless, the standards which allow this flexibility are objective both in the sense of being insensitive to opinions and convictions (in the appropriate ways) and in being common for all human beings.

Furthermore, to say as Ruse & Wilson (1986) do, that moral objectivity is an *illusion* foisted on us by our genes, strikes me either as confused or disingenuous in this context. Richard Dawkins has been arguing for some time (e.g. 1975) that reality is also an illusion foisted on us by our genes, or more fashionably, that our everyday reality is really "virtual reality." (1995) Now, aside from the helpful reminder of the foundational problem of epistemology (that we have no direct access to the things-in-themselves), this is neither news, nor is it an epistemologically helpful way to put things. True though it may be that a naive realist may construe the term "illusion" to refer to cases in which we do not directly apprehend the object of our consideration, this is surely not how we ordinarily use the term. The difference between true perception and illusion is just the difference between the proper functioning of our perceptual systems and certain characteristic kinds of malfunction. For the familiar optical illusions, the malfunction involved is not a mistake in our response to stimuli, but the fact that the environment is causing perceptions in us even though the world is not the way it is supposed to be when we have the perception in

question. Perhaps we use the term illusion in other ways as well. The point is, if the term "illusion" is not to be vacuous, then there is no reason to believe that moral objectivity is an illusion just because we are genetically "programmed" to believe that there are objective moral standards. To put it in comparative terms, the mere fact of genetic determination raises no more doubts for our moral convictions than it does for our belief in the existence of an external world, for our commitment to the law of non-contradiction, and so forth.

Teleosemantics and Moral Truth

The main problem I want to address here is the general lack of awareness of the basic mechanisms by which adapted biological functions can generate objective semantic maps between signals and states of the world, which is to say, by which they can generate objective truth conditions. So let me say a little bit about how such maps get generated in general, and how and in what conditions such maps might exist for moral language.²

² For more detailed treatment, especially regarding human indicative language/belief and animal signaling systems, see Wright (1973), Dretske (1986, 1995), Millikan (1984, 1993), Godfrey-Smith (1994, 1996), Harms (1997), Neander (1991), Skyrms (1996), and Davies (1994). Lots of work remains to be done, of course, but progress has been substantial.

A concrete example will provide analogical cases suitable to the argument presented here. Cheney and Seyfarth (1990) describe a signaling system in vervet monkeys in Kenya. The vervet monkeys have to deal with three kinds of predators, eagles, snakes, and leopards. There are specific strategies appropriate to each kind of predator. The thing to do for eagles is to look up into the sky. The thing to do for snakes is to stand up and look around on the ground. The thing to do for leopards is to run up a tree. Predictably, the vervet sentries issue three different danger signals, depending on the kind of threat perceived. The rest of the troop responds appropriately to the signals issued. The coordination required between threats, signals sent, and evasive responses upon hearing the danger signal establishes the truth conditions for the signals.³ If we let signal ‘A’ be the usual signal sent when a leopard is detected and which elicits the appropriate evasive maneuvers, then ‘A’ is true just in case there is actually a leopard there. A token of ‘A’ is false just in case there is no leopard there.

Notice that truth conditions are determined directly by the historical conditions for correct functioning of the signaling system. Having truth conditions does not depend on the prior existence of some sort of propositional meaning. When the sentry cries "A", does this have the same content as "there is a leopard here now!"? Or is it an imperative, like, "everyone run up a tree now!"? Probably neither. The vervet signals seem to occupy some middle ground between our familiar indicative and imperative modes of speech. The point is that truth in signaling systems is prior to indicative content. If you prefer to reserve the term “truth” for the correspondence of propositions to the world, I won’t quibble about the stipulation. What matters is that there are objective conditions for the correctness of the issuance of a particular signal at a particular time

³ See Skyrms 1996, ch 5, for a game theoretic analysis of this kind of signaling system.

and place, and that they rely on objective semantic maps.

Suppose, now, that our vervets have another signaling system. Let's suppose that troops are territorial and are supposed to attack members of other troops. However, since the other troops are composed of members of the same species, individuals can make mistakes about whether a given vervet is a member of its own troop or some other. So when a vervet accidentally attacks a member of its own troop, other members of the troop screech at it, causing it to desist. Supposing that this system of behavioral control is an adaptation, then it seems that there is a matter of fact as to whether the system is doing its job (in the historically adaptive manner), which is to say, whether or not the particular screech tokens are true. They are true just in case the vervet being attacked is a member of the same troop. They are false just in case it is not. Notice that if this behavioral control system is common property of all vervets, then there is an objective set of standards governing the correct issuance of these particular screeches, and these standards are the same for all troops of vervet monkeys. Whether or not the screech is correctly issued for some particular victim of course depends on in which troop it is being attacked.

Suppose, again, that vervet females get confused and pick up and carry the offspring of some other. Suppose that the adapted response is for the mother to hiss loudly at the offender, which gets the offspring returned. If this response is there in order to keep mothers and offspring sorted out correctly, then, again, there are objective truth conditions for this kind of hiss. The hiss is true just in case the female has picked up the offspring of some other. It is false just in case she has not.

If truth in simple signaling systems were all there is to biofunctional semantics, the literature would not be as large nor as difficult as it is. One of the reasons that the teleosemantic literature on meaning *is* so difficult is that what one needs to do is to work up in complexity from these simple kinds of signaling systems to ones with the characteristics of human indicative

language. Human language is fantastically sophisticated in biological terms, and a theory of meaning that hopes to account for the way truth and content work in such a system requires a lot more moving parts. Thankfully, all that is not necessary for my argument here. All that is necessary is to realize that, in principle, any adapted signaling system can have objective truth conditions for the issuance of signals, and that this does not depend on the possibility of assigning descriptive content to the signals in question. What it does depend on is whether the signal fulfills or does not fulfill its historical role in the function for which the issuing system was selected.

Consequently, human moral imperatives and claims of justification can have objective truth or falsity even though they express no descriptive/propositional content in the way that factual claims do. Of course, if you stipulate that truth requires propositional content, then this won't qualify. But remember, what is at issue is not whether or not moral language can be translated into scientific language (i.e. whether it expresses *propositions*), but whether there exist objective standards for the correct issuance of moral claims and imperatives, and whether these conditions involve semantic maps. Teleosemantics allows this without propositional content. Notice that this does not require a *separate* realm of moral facts inaccessible to scientific enquiry. All that it requires is that morality is an adaptation, and that the way that the moral hardware has paid its bills is by getting certain sorts of signals sent and responses made in particular kinds of circumstances. The signal is true if those circumstances obtain, false if not. My statement/imperative "you shouldn't do that" is morally true if the issuance of this fulfills at least my end of the function of the moral regulatory system. My claim that killing an intruder is justified because he posed a threat to my family is true just in case the language of justification is fulfilling *its* distinctive function.⁴ I submit that one needs no more than this to justify a realist attitude toward morality.

⁴ See objection 8.

Objections?

Of course, you will have a number of worries at this point. I will address the most predictable of them here.

1. *Conflicts*: The first thing one usually gets is examples which purport to show that moral norms cannot be biological norms, since morality often requires one to violate biological norms. The prevalence of this sort of objection is rather baffling, since it only takes a moment's reflection to realize that biofunctional norms conflict all the time. My eyes are supposed to see. My eyelids are supposed to protect my eyes. I cannot see when my eyes are protected, and I cannot protect my eyes when I am seeing.

2. *A similar class of objections* has the form: "just because I have X in order to Y, this does not make it moral for me to Y with X." To be sure, but who ever said that was the case? The proposal is that moral truth derives from the historical conditions for fulfillment of a specific function. There are other functions which have nothing to do with the function of morality, and thus whose fulfillment is orthogonal to that of morality.

3. *Cultural Relativism*: there does seem to be an extent to which moral truths are properly relative to cultures. If one is morally required to observe local laws and conventions, then this cannot be based on biological function, since the genetic determination is invariant. Moral truth seems to be determined by cultures, not genes.

The response is that it is well within the logic of adaptation to have flexible responses such that we are supposed to respond to local conventions. In such cases it may be that fulfillment of a biological function requires our allegiance to cultural norms.

4. *Disquotation*: I seem to be claiming that naturalism can specify the truth conditions for moral utterances, so science would seem to be able to say "'M' is true", where 'M' is some moral

claim. Isn't saying that "'M' is true" just the same as saying 'M' truly?

No. Upon observing a vervet sentry crying 'A', I may observe that a leopard is indeed present, and determine that the cry is true. But this is not to make the cry. Indeed, not being a vervet monkey I cannot *truly* make the cry, since I am not part of that signaling system. There is no biological function that I have that involves warning vervet monkeys of leopards. Similarly, if descriptive and moral language are part of two different signaling systems carried by human beings, there is no reason to assume that fulfilling the function of one entails fulfilling the function of the other.

Of course, human signaling behavior is largely integrated, but it is quite notable that integrating moral and factual knowledge has proven difficult. There is nothing mysterious about this. Disquotatation fails between signaling systems.

5. *Why think that evolution prefers truth?* Since falsehoods can be adaptive in particular circumstances, then truth just is not the same as fitness enhancement.⁵ This is all the more the case for moral truths, which may often require an individual to sacrifice reproductive success for the greater good.

Again, the proposal is that moral truth derives from the fulfillment of the historical conditions for the satisfaction of a particular function of a particular signaling system. There is more than one way to exploit a given signaling system for gains in reproductive fitness and not all of these are consistent with the particular historical function which sets the truth conditions for moral language. The important point is that semantics are determined by a particular function or range of functions. The fact that *other* functions (e.g. deception) can be fulfilled parasitically on the original semantics does not change the semantic rules, since the original function is not modified. There is a large literature that indicates that group adaptations in general have

⁵ Putnam 1983 makes this sort of argument.

problems with cheating. To say that cheating on a cooperative strategy is adaptive is not to say that it is not recognizable as cheating from a biological point of view.

6. *What reason do we have to think that any of this is true?* I have a two part response to this. The evasion first. I need not claim that any of this is true. My argument is that if you believe that morality is an adaptation and you want to be able to talk like a moral realist, then the resources provided by the biofunctional analysis of signaling systems to ground objective (though species specific) truth conditions are both available and attractive. The standard worries are easily addressed and an appropriate degree of cultural relativism is allowed.

The direct answer is that this is an empirical matter. We will probably never be absolutely certain whether this is true or not, any more than we will ever be absolutely certain that we are the products of four billion years of evolution on this planet. Such is the human condition. But, there are good reasons for thinking that morality is in fact an adaptation, ranging from the results of evolutionary game theory⁶ to evolutionary psychology⁷ to cultural anthropology⁸, and the new group selection⁹. Morality seems to regulate behavior in a way that is advantageous, at least at the group level, and we are beginning to understand that there are a variety of mechanisms which can overcome intra-group competition and can select for group level adaptations. It is also costly enough and distinctive enough that it seems implausible that it is selectively neutral or a mere side effect of some other adaptation.

7. *How can the mere presence of objective standards offer an argument for moral realism*, when anti-realists like Simon Blackburn and Alan Gibbard have their own version of

⁶ Axelrod 1984, Batali and Kitcher 1995, Skyrms 1996, Trivers 1971.

⁷ Barkow, Cosmides, and Tooby 1992; Wright 1994.

⁸ Sperber 1996.

⁹ Wilson and Sober 1994.

objective standards for “warranted assertability” and the like.¹⁰

To begin with, it seems to me that including objective semantic maps in establishing assertability conditions takes one beyond anything that might reasonably be called “anti-realism”, but I don’t consider the mere name worth fighting over. The important point is that the semantic maps are not something that are part of the resources of current anti-realism. The deeper point, which I have alluded to already, is that the current distinction between realism and anti-realism in ethics, or more precisely, the cognitivism/non-cognitivism distinction, rests on an assumption whose falsity should be apparent at this point. Cognitivist realism assumes that in order for there to be truth, there must be propositional content. Such propositions either predicate natural properties of natural objects (making moral truth reducible to scientific fact), or predicate mysterious non-natural properties. This is the cognitivist dilemma, both horns of which the anti-realist rejects. The foregoing analysis, in showing how truth need not depend on propositional content, shows a middle way for the realist — a way to reject both horns of the dilemma without rejecting the possibility of objective moral truths. Inasmuch as this removes the non-cognitivist’s reason for rejecting realism, it changes considerably the landscape of the realism/anti-realism debate in ethics.

8. *The role of reason.* The purpose of this paper is to show how the correspondence semantics of simple signaling systems undercuts the prevailing skepticism with respect to moral realism among naturalists. The focus on semantics allows one to avoid excessive speculation concerning the actual history and functional architecture of various aspects of moral deliberation, while demonstrating the realist potential for histories and architectures that might be empirically

¹⁰ Thanks to Eric Cave for this point.

substantiated. This focus on basic structures makes it harder to see how certain important phenomena might get explained. In particular, it is not immediately clear how the model of multiple independent control systems can accommodate the way reason and morality interact, this being one of the central problems for metaethicists.

Here is one story that might be told: signaling/control systems that evolve independently may come to be selected for interaction, much as independently evolved species may form symbioses. Reason, whatever its ancestral function, has come to be a control system which has endogenous norms (e.g., consistency) but no endogenous ends. Rather, it adopts the ends of other control systems, and makes its contributions to reproductive fitness via innovative strategies for the attainment of those ends. Reason does not identify the ends directly, of course, but rather, the imperative signals are accepted as motivating, and satisfaction of the end is determined upon cessation of the motivating signal. Reason adopts the ends of the feeding/foraging system, for instance, by being motivated by the hunger signal and devising strategies to bring about the cessation of the hunger signal. The hunger signal has its own semantics, as well as an ancestral behavioral repertoire. The reason that this arrangement has been selected is that while sometimes the ancestral feeding strategies work, sometimes they don't, and it is only some more creative strategy that stops the hunger signal. Reason adopts the ends of morality by being motivated by the internal versions of moral imperatives, and perhaps comes to represent the object of the moral impulse in reason's own propositional framework as a strategy for furthering those ends. A moral impulse demands rescue of a drowning child, which ordinarily would launch one into the water. Reason steps in and points out that one can't swim and will fail to save the child and drown oneself in the bargain. One rushes off in search of help and the child drowns. Jumping in and drowning would have been approved of as heroic — a properly functioning moral system. (Presumably there are signaling systems which exist to make such judgements.) Rushing off in

search of help is also morally justified, but *only* if (1) the action was motivated by the moral impulse to help the child, and (2) that this was *rationally* a better way to go about helping the child. In short, rationality is *morally* justified in intervening in the operation of the moral control system if the action is of the sort that the cooperative arrangement between morality and rationality has been selected to bring about. The reason that this is what constitutes justification is that the truth of claims of justification derive from another signaling system whose function is to evaluate whether or not this sort of creative action does in fact fulfill the appropriate requirements.

Whatever the real truth of the matter, it is after all not so hard to see how reason and morality might interact. Reason exists to maximize (or at least increase) the satisfaction of desires. There are usually more direct ancestral strategies for satisfying desires, but these don't always work, especially in changing environments. The desire to help the child is a legitimate moral desire if it is generated by the correct functioning of the moral system. It is also a rational desire if it is the kind of desire that rationality has been (recently) selected for satisfying. Rational interference in "reflexive" moral behavior is justified if it is motivated by the desire and does a better job than the reflex at satisfying the desire.

Conclusion

The full title of Millikan's 1984 book is *Language, Thought, and other Biological Categories: New Foundations for Realism*. The main thrust of this paper is to point out that, due to the way her biofunctional semantics builds up to indicative/propositional language from the simpler representational systems of our simpler brethren, whatever realism can be supported this way works equally well for moral as for scientific language. Furthermore, the possibility that we have multiple (quasi-) independent adapted signaling systems may have important implications for

issues of the reducibility of moral norms to natural science.

Some will think that the realism offered is not realist enough, for either kind of language. On Millikan's account of indicative language, one generally does not have reflective access to the meaning of one's thoughts¹¹, and we find in extending to the moral case that whatever standards exist are specific to the family of organisms carrying the moral system. Such standards are clearly not as universal and eternal as many have desired. But I fail to see why standards rooted in our adaptive history (which is, after all, fixed for all) are not objective enough to justify confidence in moral realism. On such an account, to be a human being means something, and requires something of us. There are rules for being human, and the fact that they derive from the historical process of our making rather than from the implications of some definition which we happen to satisfy need make them no less binding.

Of course, whether or not the historical facts turn out to be such as to account fully for our moral life, putting objective standards where we think they are and flexibility where we expect it, is an open question. I am optimistic about this, but whatever the case, it is time for naturalists to stop apologizing for the fact that the evolutionary view of human nature makes a sham of morality, or leaves it an utter mystery. It most certainly need not, and it probably does not. On the contrary, biofunctional semantics offers a robust grounding for moral realism in the natural world.

References

- Axelrod, Robert. 1984. *The Evolution of Cooperatino*. New York: Basic Books.
- Barkow, Jerome, John Tooby, and Leda Cosmides. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.
- Batali, John and Philip Kitcher. 1995. "Evolution of Altruism in Optional and Compulsory

¹¹ This is the rejection of "meaning rationalism".

- Games." *Journal of Theoretical Biology* 175:161-71.
- Cheney, D. L. and R. M. Seyfarth. 1990. *How Monkeys See the World*. Chicago: University of Chicago Press.
- Davies, Paul Sheldon. 1994. "Troubles for Proper Functions." *Nous* 28(3):363-81.
- Dawkins, Richard. 1989<1976>. *The Selfish Gene, New Edition*. Oxford: Oxford University Press.
- _____. 1995. "The Evolved Imagination: Are Our Brains Hard-Wired for Virtual Reality?" *Natural History* 104(9):8.
- Dretske, Fred. 1986. "Misinformation." In *Belief: Form, Content and Function*, edited by R. Bogdan. Oxford: Clarendon Press.
- _____. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Godfrey-Smith, Peter. 1994. "A Modern History Theory of Functions." *Nous* 28(3):344-62.
- _____. 1996. *Complexity and the Function of Mind in Nature*. Cambridge University Press.
- Harms, William F. 1997. "Reliability and Novelty: Information Gain in Multi-Level Selection Systems." *Erkenntnis* 46:335-63.
- Kitcher, Philip. 1993. "Four Ways of "Biologizing" Ethics." In *Evolution und Ethik*, edited by K. Bayertz. Reclam.
- Millikan, Ruth G. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. A Bradford Book. Cambridge, MA: MIT Press.
- _____. 1993. *White Queen Psychology and Other Essays for Alice*. A Bradford Book. Cambridge, MA: MIT Press.
- Neander, Karen. 1991. "The Teleological Notion of Function." *Australasian Journal of Philosophy* 69(4):454-68.
- Putnam, Hilary. 1983. "Why Reason Can't Be Naturalized." Pp. 229-47 in *Realism and Reason*. Philosophical Papers <3>. Cambridge University Press.
- Ruse, Michael and Edward O. Wilson. 1986. "Moral Philosophy as Applied Science." *Philosophy* 61:173-92.
- Singer, Peter. 1975. *Animal Liberation: A New Ethics for Our Treatment of Animals*. New York: Random House.
- Skyrms, Brian. 1996. *Evolution of the Social Contract*. Cambridge University Press.
- Sober, Elliott. 1993. *Philosophy of Biology*. Westview Press.

- Sperber, Dan. 1996. *Explaining Culture: A Naturalistic Approach*. Cambridge, MA: Blackwell.
- Tarski, Alfred. 1944. "The Semantic Conception of Truth and the Foundations of Semantics." *Journal of Philosophy and Phenomenological Research* 4:341-75.
- Trivers, Robert L. 1971. "The Evolution of Reciprocal Altruism." *The Quarterly Review of Biology* 46(3):35-57.
- Williams, George C. 1994. "A Sociobiological Expansion of *Evolution and Ethics*." Pp. 179-214 in *Evolution and Ethics*, edited by J. Paradis and G.C. Williams. Princeton, NJ: Princeton University Press.
- Wilson, D. S. and Elliott Sober. 1994. "Reintroducing Group Selection to the Human Behavioral Sciences." *Behavioral and Brain Sciences* 17:585-654.
- Wright, Larry. 1973. "Functions." *The Philosophical Review* 82:139-68.
- Wright, Robert. 1994. *The Moral Animal: The New Science of Evolutionary Psychology*. New York: Pantheon Books.