

Intentionality Without Evolution: The meaning of life and more

(for *Southwest Philosophical Studies*, Vol. 28: 2006)

Normativity is certainly a central feature, if not the true hallmark, of intentional phenomena. Believers are *essentially* things that can be *correct or mistaken* with respect to the way things are. So nothing warrants being called an *account* of belief without an attendant account of error. Accounts of intentionality thus ought to provide stories (couched in suitably informative or “naturalistic” terms) about how creatures can become *beholden to* (or correct or mistaken with respect to) the way things are. This normative perspective might seem to favor so-called “teleo-biological” accounts of intentionality, insofar as they exploit the intuition that a creature makes some sort of *mistake* when it does something that isn’t “good” for it, where goodness is spelled out somehow in terms of a creature’s ability to survive to reproduce.¹ Since such accounts seek to reduce our own intentional capacities to those of simpler creatures, they tend to deny that there is any intelligible notion of “original” intentionality above and beyond the “second-class” or “derived” intentionality commonly attributed to simpler intentional systems.² In the face of persistent criticism to the effect that these theories leave the attribution of intentional states too indeterminate, too *ad hoc*, or too historically contingent to count as a story about bona-fide mental capacities,³ advocates have countered: “Well, how else could it go?” What, other than evolution by natural selection, could possibly provide us with the requisite normative oomph, for an account to be recognizable as an account of intentionality? My aim in this paper is to respond directly to this challenge by profiling a type of beholdenness to the way things are that is intelligible as such without any blatant appeals to biological norms or to what Dennett calls “Darwin’s Dangerous Idea.”

¹ See, for instance, Millikan [1984], [1993] and Papineau [1993], [1998].

² As Dennett puts the point, “My view is that belief and desire are like froggy belief and desire all the way up. We human beings are only the most prodigious intentional systems on the planet, and the huge psychological differences between us and the frogs are ill described by the proposed contrast between literal and metaphorical belief ascription.” (*The Intentional Stance*, p. 112)

³ Fodor rehearses several of these complaints in his [1990].

Let's begin with Dennett's thesis that the attribution of intentional states is appropriate for and only for those beings whose behavior falls into *rational patterns* discernible from an intentional stance.⁴ Dennett has never been entirely forthcoming about what constitutes a rational pattern of behavior. In fact, his recent attempts to debunk the notion of original intentionality are of a piece with a failure to consider the possibility that there might be *several varieties* of such patterns corresponding to different *types* of intentionality. If however we reject Dennett's reasons for abandoning the quest for "real meaning,"⁵ we are in a position to look for different *kinds* of rational patterns corresponding to different ways in which one can adopt an intentional stance.

So what would make a pattern of activity rational? Let's take a glimpse at the poster-child of the biological approach to intentionality, the frog of philosophical legend, who famously stuffs itself silly with lead pellets. The frog isn't a compelling example of genuine mental capacity because it seemingly fails to respond rationally to its mistakes, or to respond to them *as mistakes*. One might say that it isn't capable of getting things right or wrong "by its own lights." This suggests that we'd want to have some account of *self-correction* in order to capture behavior appropriately governed by an *acknowledgement* of the norm of correctness. But in order to sustain the claim that such activity is appropriately regarded as involving the *correction* of *errors*, it would seem that we must have some account of a subject's aims as well, for how could we recognize mistakes *as mistakes* unless they are somehow liable to prevent a subject from attaining its desired ends? That is, it would be difficult to tell a story with the requisite normative punch without including some account of a creature's goals. It is thus reasonable to

⁴ See in particular Dennett [1987], ch. 2.

⁵ And we should! The demand for an account of intentionality that isn't biologically grounded – one might say "*original*" – is not as unintelligible as some advocates of the biological approach would have us suppose. In particular, we should not be persuaded by the argument that since we are products of natural selection, the ways in which we can be understood as correct or mistaken with respect to the way things are are intelligible *only* with respect to Mother Nature's standards or intentions. (See Dennett [1987], ch. 8 and [1995], ch 14.) While biological accounts of intentionality might succeed in showing how we *can* attribute to humans intentional states like those of simpler organisms, they don't thereby show that this is the *only* type of intentionality properly attributed to humans. The fact that we can be understood as having biological intentionality doesn't *preclude* the possibility of our possessing non-biological types of intentionality *as well*. Having biological intentionality doesn't preclude the possibility of there being yet *other* ways to evaluate human behavior that turn out to be largely independent (or intelligible apart from) the purposes for which humans have been designed or selected. For more discussion, see Beisecker [2002].

suppose that discernibly rational activity requires elements of both critical (self-corrective) and practical (means-end) reasoning. So to a first approximation, I propose we regard a pattern of activity as discernibly rational if it exhibits self-corrective behavior that is directed towards some goal.

I would argue that the activity of certain educable creatures is discernibly rational in precisely this sense. I admit this claim is far from novel; instead, what is new is the route by which I defend this thesis. While it is widely maintained that creatures with educable capacities enjoy richer mental lives than hardwired, tropistic beings,⁶ accounts of educable capacities typically dwell upon how they render organisms better able to fulfill their natural purposes in the face of environmental contingency. Dretske, for instance, focuses on how providing creatures with the ability to conduct their own selection of appropriate internal indicators might be the best way for a designer (including Mother Nature) to solve the problem of constructing creatures that are likely to fulfill their intended purposes.⁷ Unable to anticipate the relevant regularities of a creature's environment, a designer might find it useful to equip creatures with some ability to tailor their own responsive dispositions to their particular surroundings. While this thought is probably correct (as far as it goes), it doesn't show how creatures with such educable capacity exhibit any *intentionality* that is intelligible as such, *apart* from the purposes for which they have been selected or designed. At best, it explains why behavioral plasticity might have been favored by selective pressures. Although there is this vague intuition that educable creatures are responsive to error, and so "learn from their mistakes," it remains to be seen how these mistakes could be intelligible as such without ultimately appealing to the creatures' biological purposes.

If we had a story about how the flexibility of educable creatures gives rise to a different, non-biological sort of accountability or susceptibility to evaluation, then we might begin to see how an *original* intentionality could be a product of natural selection. My first task is to describe a specific type of behavioral plasticity, which I will call

⁶ In particular, see Dretske [1988], [1999]; Bennett [1976], [1991]; and Dennett [1995].

⁷ See Dretske's discussion of the so-called "design problem" in his [1988], pp. 96ff.

“expectation-mongering,” that is free of any appeals to survival value or selective purposes. Here’s how the story goes. To account for certain “blocking effects,” several learning theorists have argued that the observed educability of some animals is best explained in terms of the adjustment of “expectation-like” structures mediating between sensory input and behavioral output.⁸ Accounts of expectation-based educability aim to make sense of (that is, to describe in suitably informative terms) patterns of activity, whereby creatures exhibit an appropriate sensitivity to the consequences of their own responses. The basic idea behind these accounts is that the actual responses such creatures make in situations is a function of various outcomes that they currently associate with the particular responses available in their behavioral repertoires. Since different responses in the same situation can bring about different outcomes, and since the same type of response can, depending upon the circumstances, yield different outcomes, the structures these accounts posit to mediate between sensory input and behavioral output – let’s call them “expectations” – need to include (at least) three separate components: 1) conditions of activation and deactivation, 2) a response type, and 3) a consequence condition. The first component specifies, as it were, when an individual expectation is turned on or off. When an expectation is activated, the creature associates the outcome specified by the third component with the response specified by the second.⁹ Should the creature engage in that response and the consequence condition

⁸ See, for instance, Staddon [1983], pp. 414ff and Dickinson [1989]. For example, some animals that have been trained to associate a conditioned stimulus with an unconditioned stimulus will subsequently fail to associate other stimuli with the unconditioned stimulus, when the latter are presented along with the original conditioned stimulus. Rats that have been trained, for instance, to associate a bell tone with an electric shock will not come to associate a red light with a shock, as long as the red light is consistently paired with the bell tone. The prior conditioning prevents (or “blocks”) subsequent conditioning to other, co-varying stimuli. If this learning were merely a matter of the frequency of stimulus-pairing, then one would expect the animal to become conditioned to the new stimulus as well. One would expect the rats eventually to associate the red light with a shock, as indeed they do when they aren't subjected to the earlier training. Many learning theorists have argued that the failure of previously conditioned animals to become conditioned to the new stimulus arises because the animal already uses the original conditioned stimulus to *predict* the occurrence of the unconditioned stimulus, and with a reasonable degree of success. When a previously conditioned rat encounters the compound tone and light stimulus, it *expects* that the shock will occur (because it heard the bell tone), and so the subsequent shock isn't a surprise. Since events are as they were *expected* to be (they were not novel), there is no pressure to develop new associations, and there is no subsequent conditioning to the light. Thus these theorists conclude that the rats are responding to *surprise*, to things not being as they *expected* them to be.

⁹ To be relentlessly naturalistic, the first and third components could be specified in terms of, say, activity along a creature’s sensory manifold, while the second in terms of the activation of particular motor programs.

not to be satisfied, then the creature would be disposed to adjust the components making up that expectation. Through the revision of individual expectations when they are so “violated” these creatures distinguish themselves from the merely tropistic, and so display the sensitivity to the consequences of their own responses that learning theorists have sought to describe.¹⁰ So the story is basically this: under certain circumstances, an expectation will be activated and the creature will then anticipate that a certain response will yield a particular outcome. Should that turn out not to be the case, its dispositions to form such anticipations will change. *Expectation-mongering creatures* can now be defined as those whose overall behavior is most systematically described as governed in part by the consequence conditions of currently activated expectations.¹¹ For example, a creature might be disposed to engage in any responses associated with the outcome of its coming into acquiring cookies, or it might be disposed *not* to engage in any response associated with electric shocks.

Given the intuition that we ought to be able to evaluate expectations as correct or mistaken, this would seem to be a promising beginning of a story about intentionality that isn’t biologically grounded. Note that the description I gave of expectation-mongering behavior doesn’t make any obvious appeal to biological purposes. One can identify expectation-mongering creatures as such without having to recognize them as subject to selective pressures. Nor have I construed expectation-based educability as the selection of responsive dispositions that have positive survival value, although that is presumably

¹⁰ Different accounts of expectation-based educability differ with respect to which expectation components are allowed to vary from expectation to expectation, which components are capable of being altered, and also the conditions in which they stand to be adjusted. Staddon [1993], for instance, takes learning to involve the adjustment of consequence conditions, while Bennett [1976] effectively restricts it to the revision of activation conditions. A fully general account of expectation would leave as much of this up for grabs as possible.

¹¹ More formally: their response to situations is a partial function of the consequence conditions of currently activated expectations. It bears mentioning that I’m not trying to show that any particular creatures are expectation-mongers. That is the work of ethologists, not philosophers. Notice also that I’ve defined expectation-mongering in terms of how a creature would behave in various possible situations. Since any pattern of actually observed behavior could be the product of tropisms, showing that a creature is an expectation-monger would have to involve establishing that certain counterfactuals hold. It turns out, then, that those who *design* devices would likely have an easier time justifying the attribution of expectations to their subjects than those who encounter them “out in the field,” simply because they have a better sense of what goes on inside the “black boxes” they are studying, and so would have a better grasp of the relevant counterfactuals. For a discussion of the difficulties attributing to wild subjects states similar to the expectations described here, see Heyes and Dickinson [1990].

something such behavioral plasticity can bring about. However, remember that to sustain the claim that expectation-mongering creatures exhibit a non-biological sort of intentionality, we need to show how their behavior fits an overall *rational pattern*. That is, we need to show how expectation-mongering can be viewed as goal-directed, self-corrective activity in its own right. Here my strategy will be to begin with the practical side, and construct an account of goals from this account of expectation-mongering, and then turn around and use this account to ground the notion of expectation error.

Fortunately, the definition of a goal turns out to be satisfyingly straightforward and intuitive: a certain outcome is to be regarded as one of a creature's current goals, to the extent that the creature is disposed to engage in responses *expected* to bring about that outcome.¹² For example, the creature mentioned above that is systematically disposed to engage in responses associated with the outcome of acquiring cookies can be understood as having the acquisition of cookies as its goal. The other creature, disposed not to engage in responses associated with an electric shock, can be understood as having an *aversion* to shocks. Goals that are construed like this work in conjunction with a creature's expectations to *explain* its particular responses to situations. By characterizing responses in terms of the outcomes they are expected (by the subject) to bring about, these explanations show how a particular response fits a creature's overall pattern of responsive dispositions. And we need not regard such explanations as empty, because they point out that a subject might have done otherwise, had that response not been expected to bring about a certain outcome, or had some other response been expected to bring about that outcome instead.¹³ Notice in particular that an expectation must have an appropriate consequence condition before it can be paired with a goal in order to explain a creature's behavior. The expectation's *content* - that is, its consequence condition - must itself satisfy the goal's *condition of satisfaction*. Since goals and expectations must have the right sort of "fit" with one another before they can successfully explain a creature's

¹² In a similar fashion, we can determine a (possibly non-transitive) preference ordering among outcomes. Observe, by the by, how goals so-construed rest upon an antecedently intelligible account of expectation, and as such, are not conceptually prior to expectations.

¹³ I am aware that certain hard-nosed physicalists (e.g., Kim) might look askance at such "dispositional" explanations of behavior as, at best, incomplete. I have not set out to assuage the concerns of physicalists,

behavior, these explanations face what could be thought of as a *rational constraint*. Thus it makes some sense to claim that attempts to explain a creature's behavior with respect to its goals and expectations to be attempts to *rationalize* its behavior.

While this is obviously a broadly dispositional account of goals, it does not crudely identify a creature's goals with the actual outcomes that the creature is likely to bring about.¹⁴ On this proposal, creatures do not have to be disposed to bring about the eventual attainment of their goals. For one thing, just as we can pick out fragile objects without requiring that they manifest their fragility by shattering, we can identify a creature's goals, even though it might not ever find itself in circumstances where their attainment is possible. For our purposes, however, the respects in which the activation of *expectations* can block the attainment of goals are particularly significant. Here we can say that an expectation-mongering creature will be disposed to attain its goals (whenever such attainment is possible) to the extent that its expectations are configured *correctly*.¹⁵ This, of course, is where the normative rabbit gets pulled out of the naturalistic hat. The nice thing is that we can pick out unfavorable expectation configurations likely to hinder a creature's attainment of its goals, and so have reason to regard these configurations as expectation *errors*. For instance, a creature is liable not to fulfill a goal if one of its expectations is activated in a situation in which the expectation's response would fail to bring about the satisfaction of its consequence condition. A creature is likely not to fulfill its goal of acquiring cookies if a response it associates with the outcome of acquiring cookies will actually bring about some different outcome. We can thus think of such an occurrence as an *error of commission*. Similarly, an *error of omission* arises whenever the response of an expectation that is not activated would bring about the satisfaction of its consequence condition (that is, were its activation not to be an error of commission). Here our creature is liable not to engage in a response that would procure cookies, since it fails to associate that response with that desired outcome. Since these two expectation

who are generally suspicious of intentional explanations anyway. There are plenty of others (e.g., Dretske) engaged in that project! In particular, see Bennett [1991] and Dretske [1988].

¹⁴ It's also worth remarking that this account of goals is neither a "reinforcement" (see Whyte [1993] and possibly Dretske [1988]) nor an "extinction" (such as that occasionally attributed to Russell) theory.

¹⁵ As Bennett might claim, a critter will be disposed to attain its goals "all things being equal," and having correctly configured expectations is part of things being equal (Bennett [1990], pp. 42ff).

configurations are liable to prevent a creature from attaining its goals, expectation-mongering creatures are susceptible of two distinct sorts of mistakes about the way things are in their environments.¹⁶ They can be evaluated as having gotten things right or wrong, and so can be understood to exhibit a type of intentionality above and beyond that typically attributed to artifacts and simple organisms. Observe once more that while expectations are, as it were, *ontologically* or conceptually prior to goals in the sense that the latter can only be defined in terms of an antecedently intelligible account of the former, goals nevertheless enjoy a *normative* priority over expectations in the sense that the notion of expectation error depends upon (or is intelligible as such only with respect to) this account of goal-directedness. A slogan: while goals owe their existence to expectations, expectations owe their normativity to goals.¹⁷

So we now have reason why, from a creature's own perspective, its expectations *ought* to be activated just in case their consequence conditions would be satisfied, were the creature to engage in the response picked out by that expectation's response component. As an account of error, this story has several appealing features. Heading that list is the fact that, unlike biological accounts, the commission of these errors doesn't depend upon any antecedent determinations of when given responses tend to have survival value or to be reproductively advantageous for a creature. In addition, these standards for expectation correctness are *categorical* – one might say “anti-pragmatic”- in the sense that they apply as they do, irrespective of the particular goals a creature might possess. The activation (or inactivation) of an expectation can be identified as correct or mistaken, regardless of what a creature's goals happen to be. Moreover, the conditions for the appropriate activation of one expectation can be quite different from the conditions of appropriate activation for another. That is, the activation of separate expectations can be beholden to distinct features of a creature's environment. As a result of this *feature selectivity*, expectation-mongering creatures can be correct with respect to some features of their environment, yet mistaken with respect to others. They can get

¹⁶ Please observe that this account doesn't rule out “accidental” (or unexpected) success at attaining goals.

¹⁷ Insofar as the goals so-construed rest upon an antecedently intelligible account of expectation, this account reverses the strategy historically advocated by Ramsey, and more recently pursued by Whyte

things right or wrong *in a variety of respects* due to the simultaneous activation of several expectations. In fact, an expectation-mongering creature could even be *massively* mistaken about the way things are.¹⁸ Furthermore, the situations in which one expectation would be appropriately activated might just happen to line-up or co-vary with those in which another would be activated. For instance, the circumstances in which one response would procure cookies might be precisely those in which another response would bring on an electric shock. “Extensionally speaking,” distinct expectations can thus share the same circumstances of appropriate activation. However, the *particular means* by which these circumstances are picked out would differ for each such expectation, simply because they would be comprised of different expectation components. So even though their circumstances of appropriate activation can be the same, their content (“intensionally speaking”) can remain quite distinct. Had the subject’s environment been otherwise, these expectations might not have shared circumstances of appropriate activation. It would thus appear that attributions of expectation states exhibit something like the ballyhooed semantic opacity or sensitivity to intensional contexts so often associated with the attribution of genuine intentional states. To attribute an expectation to a creature is not tantamount to attributing to it other expectations sharing the same circumstances of appropriate activation.¹⁹

In addition to the practical rationality I’ve just described, expectation-mongering creatures can clearly exhibit a certain measure self-corrective, *critical* rationality as well. Insofar as they are disposed to revise their expectations in the wake of the errors described above, such educable creatures take discernibly rational steps to minimize future mistakes. There is of course no guarantee that these revisions will yield future

[1993] and Papineau [1998]. I’m not sure exactly what to make of it, but I find it intriguing that normative and ontological orders of explanation can come apart from one another.

¹⁸That is, this story does not appeal to so-called “normative constraints” on the attribution of intentionality, such as Dennett’s “assumption of rationality” or Davidson’s “principle of charity”; interpreters are not constrained to attribute expectations that are for the most part correct. See Dennett [1987], pp. 17ff. and Davidson [1984], p. 196.

¹⁹ It would seem then that we’ve found grounds to challenge the popular contention that a fine-grained sensitivity to intensional contexts would require linguistic capacity. See, for instance, Davidson [1985], pp. 474-476. This account thus meets what might be called “Davidson’s challenge.” At the very least, advocates of Davidson’s position would need to clarify just what they mean by the sensitivity to intensional contexts alleged to be required for legitimate ascription of intentional states.

success.²⁰ The point is just that creatures displaying this sort of educable capacity would take expectation correctness or aptness to be a *regulative ideal*, at least in the sense that they are disposed to revise error-prone expectations while leaving correct expectations as they are. And so it seems that they display something akin to rational responsiveness to error that Davidson argues must be possessed by any rational animal.²¹ By responding in a more or less reasonable fashion when the outcomes of their responses aren't as they were expected to be, such creatures manifest an apparent capacity to be "surprised."

In sum, we can discern a rational structure in the behavior of expectation-mongering creatures. As I've shown, they can pursue goals based upon possibly mistaken ideas about how to attain them, and their ability to attain these goals can improve through experience. Since they can be evaluated as having gotten things right or wrong with respect to the way things are, we are justified in crediting these creatures with a kind of intentional capacity that avoids any obvious appeals to natural purposes or proper functioning. Indeed, since expectation-mongering creatures wouldn't have to be products of any sort of selection, natural or otherwise, and their expectation errors are intelligible as such without our having to consider the purposes for which they have, as it were, been designed, this account shows how non-biological "creatures"- for instance, those philosophical fantasies spontaneously generated out of swamp muck - could nevertheless possess a certain kind of intentional capacity. Thus the account avoids the awkward conclusion that a being physically indistinguishable from something capable of bonafide mental representation could nevertheless be wholly incapable of getting things wrong, on account of its lacking a suitable, biological pedigree.²² Moreover, not only is this account of goals intelligible apart from considerations of a creature's biological purposes, these goals might even collide with those purposes. For instance, there is no

²⁰ Against the background of this account of expectation error, we can understand educable creatures to be making *errors of expectation revision* whenever they adjust an expectation in ways that would render it more susceptible to either errors of commission or errors of omission.

²¹ See Davidson [1984], [1985], and [1999]. To be sure, Davidson tries to argue that the conceptual resources required to be surprised in turn require an animal to be capable of interpreting the utterances of others; thought requires talk. However, we don't have to accept this argument to take Davidson's point that the capacity to be surprised, or to recognize when the way things are aren't as one took them to be, is an important part of being a rational animal.

²² Of course, such a being would lack an appropriate history to possess biologically grounded sorts of intentionality.

reason why a creature couldn't be disposed to respond in ways expected to bring about reproductively disadvantageous outcomes. Such a creature would have a goal that is, from a biological point of view, remarkably maladaptive.

3. Ecumenical Conclusion

In this paper, I've shown how a certain type of educable creature can be correct or mistaken with respect to the way things are in a way that isn't grounded in the functions or purposes for which they've been designed or selected. Such creatures thus may be said to exhibit an *original* type of intentionality above and beyond whatever biological intentionality they might possess. So I think we have a constructive refutation of the claim that *all intentionality* is intelligible as such, only by appealing to natural purposes. Yet this conclusion remains pleasantly ecumenical in the sense that it recognizes that there can be several distinct kinds of intentionality, including that adumbrated by biological accounts. Different accounts of intentionality need not compete with one another. This ecumenicism requires me to end on a note of caution. My task here has simply been to show that it makes sense to talk about intentionality that isn't biologically grounded. In particular, I'm not about to claim that the admittedly simple type of non-biological intentionality I've just described accounts for the more sophisticated intentional capacities exhibited by beings like us. We linguistic creatures evidently engage in performances that can be evaluated as correct or mistaken with respect to the way things are according to standards that are instituted across our linguistic communities. I would agree with Davidson that this story about expectations is far too atomistic and individualistic to capture our apparent capacities to justify our worldly commitments to one another. Nevertheless, the story I've already told could prove to be a good *platform* upon which to erect *further* accounts that set these more involved intentional capacities as *targets*. Indeed, the fact that we can describe in broadly naturalistic terms a pattern of activity in which an (albeit primitive) *original* type of intentionality can be discerned should provide us with some hope that we can describe other patterns in which more sophisticated kinds of intentionality can be described,

including the irreducibly social stripe that we enjoy. I have some thoughts about how such stories could go – thoughts which draw from the broadly pragmatist idea that the meaning of a sign is some function of its *expected* consequences – but I’m afraid such ruminations will have to await another time.

Works Cited:

Beisecker, David (2002) "Dennett and the Quest for Real Meaning: In defense of a 'Myth'," *Philosophy in the Contemporary World*, Volume 9:1 (Spring-Summer), pp. 11-18.

Bennett, Jonathan (1976). Linguistic Behaviour. Cambridge: Cambridge. Reprinted 1990 by Hackett, Indianapolis, IN.

Bennett, Jonathan (1991). "Folk-psychological explanations," in Greenwood (ed.), The Future of Folk Psychology: Intentionality and Cognitive Science. Cambridge: Cambridge: 176-195.

Davidson, Donald (1984), "On the Very Idea of a Conceptual Scheme," in Inquiries into Truth & Interpretation. Oxford: Oxford.

Davidson, Donald (1985). "Rational Animals," in LePore and McLaughlin (eds.), Actions and Events. Oxford: Basil Blackwell.

Davidson, Donald (1999). "The Emergence of Thought," *Erkenntnis*, 51: 7-17.

Dennett, Daniel (1987). The Intentional Stance. Cambridge, MA: MIT Bradford.

Dennett, Daniel (1995). Darwin's Dangerous Idea. New York: Simon & Schuster.

Dickinson, Anthony (1989). "Expectancy Theory in Animal Conditioning," in Klein and Mowrer (eds.), Contemporary Learning Theories. Hillsdale, NJ: Lawrence Erlbaum.

Dretske, Fred (1988). Explaining Behavior. Cambridge, MA: MIT Bradford.

Dretske, Fred (1999). "Machines, Plants, and Animals: The Origins of Agency," *Erkenntnis*, 51: 19-31.

Fodor, Jerry (1990). A Theory of Content and Other Essays. Cambridge, MA: MIT Bradford.

Heyes, Cecilia and Anthony Dickinson (1990). "The Intentionality of Animal Action," *Mind & Language*, 5: 87-104.

Lettvin, et al. (1959). "What the Frog's Eye Tells the Frog's Brain," *Proceedings of the Institute of Radio Engineers*: pp. 1940-51.

Millikan, Ruth (1984). Language, Thought, and Other Biological Categories. Cambridge, MA: MIT Bradford.

Millikan, Ruth (1993). White Queen psychology and Other Essays For Alice. Cambridge, MA: MIT Bradford.

Papineau, David (1993). Philosophical Naturalism. Oxford, UK: Blackwell.

Papineau, David (1998). "Teleosemantics and Indeterminacy," *Australasian Journal of Philosophy*, 76:1, pp. 1-14.

Staddon, J. E. R. (1983). Adaptive Behavior and Learning. Cambridge: Cambridge.

Whyte, J. T. (1993). "Purpose and Content", *British Journal for the Philosophy of Science*, 44: 45-60.