Henry Jackman
Presented at the May 2000 Meeting of the Canadian Philosophical Association.

Conventionalism, Objectivity and Constitution

Most philosophers, whether they consider themselves naturalists or not, typically think that we can talk and think about objects in the world which are themselves independent of our conception of them. Unfortunately, providing a naturalistic explanation of how we can do this has proved notoriously difficult. In particular, naturalistic accounts of intentionality have had problems explaining how we can (collectively) *misidentify* objects in the world. Without the possibility of such misidentification, our intuitive sense of objectivity (and hence of objects) seems compromised. In his recent book, *Having Thought*,[1] John Haugeland provides a two pronged account of our interactions with the world that he takes to secure a degree of objectivity that many current naturalistic accounts lack. Haugeland's analysis of intentionality and objectivity in terms of the interplay of two types of 'recognitional' skill represents an advance over many standard naturalistic accounts. Nevertheless, it will be argued here that his inegalitarian conception of the two sorts of skill leaves him with a quasi-conventionalist account of our relation to the world. Consequently, his account still lacks the more robust sort of objectivity that a more holistic development of his theory could provide.

Haugeland thinks that there is a sense in which objectivity is tied to "constitution,"[2] and perhaps the paradigmatic example of a constituted domain is the game of chess. For instance, knowledge of the constitutive standards[3] of chess allows us to recognize the rooks, queens and pawns that we interact with as the objects that they are. While there are certainly important differences between our interaction with chess pieces and our interactions with other empirical phenomena, the 'game' metaphor will be used by Haugeland throughout. In any case, according to Haugeland, for each constituted domain there will be associated two sorts of skills.

---

[1] Cambridge: Harvard University Press, 1998. The book is a collection of essays, though most of the material discussed here is from his 'capstone' paper "Truth and Rule Following", which is new with the collection, all references, unless stated otherwise will be to this book.

[2] I can't do full justice to the intricacy and detail of Haugeland's position in this brief format, but what follows should capture the essentials of Haugeland's account.

[3] Which include more than just the constitutive rules (see pp. 320-1).

1. *Mundane Skills*: "the resilient abilities to recognize, manipulate, and otherwise cope with the phenomena of the game, including the other players, as required and permitted by the rules -- in effect, the ability to engage in play." (323). In the case of chess these would include the ability to recognize and move the pieces, to tell what counts as moving in a diagonal rather than a straight line, that one's king is in check, etc.
2. *Constitutive Skills*: the "resilient ability to tell whether the phenomena governed by some constitutive standard are, in fact, in accord with that standard" (323). These involve the ability to understand which moves and combinations of moves recognizable with the mundane skills are, or are not, in accord with the constitutive standards.[4]

Haugeland's project is to "show how some mundane skills, understood in this sense as interdependent with constitutive skills, can therefore be understood as *objective* -- that is, having independent objects that are criterial for their correct exercise" (325), and doing so will involve the *interplay* between the mundane and constitutive skills.

From such a perspective, various naturalistic accounts of intentionality can't generate a sufficiently robust notion of objectivity precisely because they tend to focus *exclusively* on the mundane skills of organisms.[5] If considered on their own, particular exercises of mundane skills can't be characterized as mistaken. If one works *only* with mundane skills, then the plurality of one's mundane findings can give one no ground to say that a currently problematic exercise is mistaken. This is the source of the 'disjunction problem' associated with so many naturalistic accounts of semantics. If intentionality is understood *solely* in terms of our recognitional capacities (mundane skills), it becomes virtually impossible to make sense of *mis*recognition. If I am disposed to 'recognize' both deer and cows-at-a-distance as deer, then why should one not

[4] These two sorts of skills are, of course, interdependent:

> Constitutive skills depend in an obvious way on mundane skills: in order to tell whether constituted phenomena are in accord with the constitutive standards, it is necessary to be able to tell what those phenomena are. . . . Mundane skills depend, in a different but equally obvious way, on constitutive skills. There could be no telling which pieces are where on the board if there were no such thing as chess pieces -- which is to say, if there were no constitutive standards applicable in practice via constitutive skills. (324)

[5] Or, as is often the case, focus exclusively on organisms that have only mundane skills. The frog that eats all small objects that fly though its field of vision being perhaps a paradigm case of such a creature.

say that by "deer" I simply mean *deer or cow-at-a-distance*? Many naturalists have tried to answer this problem in terms of privileging a subset of the mundane skills,[6] but there is a growing consensus that such theories are unable to makes sense of certain types of error we intuitively take to be possible.[7]

It is this failure of an individual's mundane skills to underwrite a sufficient degree of objectivity that has led some to understand normativity and objectivity in *social* terms.[8] It is then the mundane skills of the *community* that determine what a term refers to, and an individual's exercise of a mundane skill can be characterized as mistaken if it is out of line with the communal norm. However, such social accounts of objectivity lack the resources to explain how *everyone* might be wrong about a topic, and as Haugeland correctly stresses, "There is a fatal flaw in any proposed account of objectivity that cannot give sense to the possibility of *everyone* being consistently wrong about something" (315). There are, of course, phenomena for which a social account of objectivity may be correct (what the group counts as a proper greeting just *is* a proper greeting, etc.). Nevertheless, we tend to think of many objects as having an independence that is more robust than that found in such 'instituted' phenomena as greetings. Haugeland attributes the comparative lack of objectivity of such 'instituted' phenomena to their relying, even if on a social level, exclusively on mundane skills. As Haugeland puts it,

> Merely instituted phenomena, such as dance steps, greetings, and the appropriate occasions for them, are what they are only in relation to consensual proprieties of performance. Thus, it can never be the case that everyone is systematically wrong about them, takes them to be other than they are; for, if everyone systematically took them some other way, then that's what they would be. Such phenomena have, therefore, no independent normative standing over against the general [mundane] skills and customs for producing, recognizing and interacting with them. And, consequently, they can never "resist" or "stand up to" those skills – they can never show proper performance to be in error, of the skills themselves to be in need of repair. (338)

For instituted phenomena, then, proper (that is, socially endorsed) performances are incapable of being in error, and this denies the purported objects of these performances the independence that objectivity requires. Haugeland, correctly, insists that our (collective) cognitive activities are often responsible to objects that are, in a robust way, *independent* of them. As he puts it:

---

[6] See, for instance, Dretske 1981, Millikan 1984, Fodor 1990.

[7] See Boghossian 1989, Fodor 1990, and Loewer 1997 for a discussion of the problems with such accounts. Haugeland discusses the problems with attempting to ground objectivity in terms of biological function on pp. 309-10.

> What *objectivity* demands… is that the 'objects' of objective tellings should have a determinacy and normative standing *independent* of the performance norms for those tellings, such that the tellings could be performed properly [that is, in a socially endorsed way], and still get their objects wrong." (314)

This is why Haugeland argues that objectivity is found only in *constituted* domains, where such 'constitution' involves the interplay of *both* sorts of skills.

The importance of what Haugeland calls "constitution" resides precisely in the fact that the constitutive skills help institute *incompatibility* relations between the results of the mundane skills, this in turn allows groups of mundane findings to 'gang up' on intuitively problematic identifications, and thus to characterize them as mistaken.

> Constituted phenomena, by contrast, *can* stand up to mundane skills. By virtue of the constitutive standards imposing constraints on *combinations* of results, individual loci of incompatibility can resist and refute particular proper (or improper) performances – show their results to be incorrect. In effect, the standards, by ruling out the bulk of conceivable combinations, bind the totality of actual results within the narrow bounds of possibility. And this binding together allows them, so to speak, to "gang up on" isolated performances whose results are incompatible with – impossible in the light of – the overwhelming majority. Figuratively, we can think of the phenomena as gaining the power to resist by "locking arms" against the skills, with the constitutive standards providing their grip or their ability to lock together. (338)

Constitutive skills, then, allow for our recognitional capacities (mundane skills) to keep each other in check and thus underwrite a more robust type of objectivity. As Haugeland puts it, "a test is a confrontation among exercises of two or more mundane skills *and* at least one constitutive skill" (334). For example, one's mundane recognition of a cow-at-a-distance as a "deer" is not, in itself, incompatible with one's mundane recognition of that same cow as a "cow" once one has approached it. As Haugeland puts it "in and of themselves, no two (or more) actual [mundane] findings are incompatible" (334). However, these two mundane finding become incompatible when combined with the constitutive recognition that, for instance, deer do not change into cows. The cow-findings thus 'gang up' on the isolated deer-finding, and the later exercise of one's mundane skill is characterized as mistaken. Constitutive skills thus "must be 'meta' or 'monitoring' skills vis a vis mundane performances, for their essential exercise is to watch out for incompatibilities among these results" (335).

This requires that one be *committed* to the constitutive skills that purport to bind one's mundane findings together. Otherwise, mundane findings could never be meaningfully

---

[8] See Kripke 1982.

questioned, since any purported incompatibility between them would just result in our giving up on the constitutive standard that linked them together.[9] One can't, for instance, simply give up the belief that deer and not cows in the face of the first set of mundane findings that suggests otherwise. If one did, one would be back in the original position of not having one's mundane findings open to any sort of testing.

Haugeland's insistence that an account of intentionality involve our general beliefs (constitutive skills) as well as our recognitional capacities (mundane skills) focuses on an important and too often neglected point. Nevertheless, Haugeland's discussion of the central role of constitutive skills has a quasi-conventionalist feel which, while perhaps appropriate to games like chess or baseball seems to undermine the robustness of the notion of objectivity he is trying to get.[10] That is to say, Haugeland's conception of the constitutive standards and the skills associated with them can be understood as "conventionalist" in that it is in many respects like Poincaré's account of how the axioms of geometry 'define' the meaning of geometric primitives. Constitutive standards are taken by Haugeland to have the status conventionalists assigned to the axioms of geometry: viewed from outside the conceptual framework, they collectively define the framework itself, but from within the framework, they express claims that are necessarily true of the objects within the framework.[11]

In particular, while he takes his project to be to "show how some mundane skills, understood in this sense as interdependent with constitutive skills, can therefore be understood as *objective*" (325), there is no corresponding sense in which the *constitutive* skills can be understood to be objective. This becomes especially clear when he argues that:

---

9 "If performances or skills are "revised" or "repaired" casually, at the first sign of trouble, then nothing is seriously excluded, and all "testing" is a farce. This is *why* the skills must be *resilient*: they must be able to stand up to one another, and hold their ground." (334).

10 Haugeland warns, "the use of games as examples is fraught with philosophical peril"(320), but he may not take his own warning seriously enough.

11 For a useful discussion of Poincaré's position, see Coffa 1991, esp. ch. 7. It has recently been argued that such appeals to implicit definitions can be separated from conventionalism itself (Boghossian 1997). Nevertheless, the argument of this paper should be unaffected by such considerations, since it turns on the problems associated with the conventionalist's appeal to such implicit definitions, not the 'non-factualism' associated with their position.

> The role of the constitutive skill(s) is distinctive and fundamental. The account of errors depends essentially on incompatibilities among mundane findings. . . . though there can be "confrontation" among several mundane findings and a constitutive finding, *there can be no incompatibility between the former and the latter.* Rather, it is at most the mundane results that are mutually incompatible, and the constitutive finding which "confronts" them by, so to speak, accusing them of that." (334-5, italics mine)

This account of testing gives the constitutive skills a non-negotiable (perhaps even quasi-analytic) character. Mundane findings can be mutually incompatible, but the option of seeing them as mutually consistent and incompatible with the constitutive skill is not open to question.

Of course the constitutive skills are open to *some* sorts of revision,[12] but Haugeland's discussions of these cases only reinforces the impression that the constitutive skills, and the beliefs associated with them, are, at bottom, not 'really' correctable.

> *To modify or repair the constitutive skills is to change the game itself*…. Such an option may not loom large for routine recreational games; but it can become central when … the primary challenge is to figure out whether there's a "game" (intelligible domain) at all…. If, in order to maintain or regain equilibrium, a number of constitutive and/or mundane skills must be modified together, the net effect is apt to seem "revolutionary"." (336, italics mine)

Haugeland here correctly point out that with objective phenomena, a central challenge is often to figure if there is a "game" and, if so, how to play it (see also p. 330). However, if this is what one is really trying to do, then one's constitutive skills should also be evaluable as right or wrong. What one takes to be the constitutive standards may not actually be the one's which govern the "game" one is trying to understand. If this is the case, however, then 'revising and repairing the constitutive skills' should be understood as trying to bring one's conception of the game into line with the game one has always been playing. No change of game need follow from repairing or revising the constitutive skills to bring them in line with the phenomena. This is, however, precisely what Haugeland denies when he claims that "to modify or repair the constitutive skills is to change the game itself." Here it seems as if our conception of the game is, by definition, correct, and our changing this conception will be changing the game. This may be a plausible line to take about constitutive skills in actual games such as chess or baseball, but it seems far less plausible when extended to our other interactions with the world.

---

[12] "To point out that constitutive skills are essentially different from mundane skills is not at all to suggest that their exercises are immune to revision, or that they themselves are immune to repair" (335). The fact that, unlike in games, both sorts of skills are subject to repair in empirical inquiry is stressed again on p. 343.

This idea that constitutive skills are 'by definition' correct, shows itself again in Haugeland's discussions of the so-called 'excluded zone':

> Every constituted domain requires an excluded zone -- a non-zero extension of the conceivable beyond the possible -- that is *in fact* empty. It is excluded in the sense that "phenomena" in it are not possible; they are ruled out by the constitutive standards. But such phenomena are nevertheless conceivable (hence, the extent of the zone is non-zero) in that, were they to occur, they could be recognized by exercises of mundane skills that belong to the domain. *Therefore*, the crucial requirement that they don't in fact occur is subject to empirical test. The availability of such tests is what lets disclosure be *hard* -- lets it be an *achievement* -- because it is vulnerable, liable to fail. (333.)

This talk about the "excluded zone" in, as it were, ontological rather than epistemic terms seems of a piece with treating the constitutive skills as not being revisable without "giving up the game."[13] If the excluded zone were understood in epistemic terms, then the discovery of an object within it could cause us to recognize that our conception of the constitutive norms was mistaken, not that the domain in question was not disclosable at all. Treating the excluded zone ontologically leads Haugeland to the type of thing he says about objects bellow:

> Constitution, making sense of objects, is not free. It depends on an equilibrium among a number of constitutive and mundane skills -- an equilibrium which, since it excludes the bulk of what it renders conceivable (testable), is empirically precarious. But, for the same reason, and for as long as it lasts, it is also an empirical achievement. The constituted objects participate in that achievement, deriving determinacy and normative status from it. But they also have an ability to resist that transcends that participation, because it is they, and they alone, upon which the equilibrium itself rests; by not cooperating, they have the power to bring it down. Even though in so doing they must annihilate themselves as what they are, nevertheless they can. (353)

This last bit about objets 'annihilating themselves' gets to the heart of what sort of objectivity seems to be missing with the constitutive skills. Since they constitute the domain in question (and are hence 'analytically' true) constitutive standards and the skills associated with them are, by definition, correct: if a 'rook' does not agree with the constitutive standards for rooks, then it ceases to be one (it 'annihilates itself'). Constitutive skills turn out not to be up for "empirical test" at all, and can thus never turn out to be mistaken: they can only turn out to be unworkable -- that is to say, our world turns out not to be 'disclosable' in terms of them.[14]

---

[13] "Constitutive commitment to chess is not an agreement to play by the rules, on pain of being rejected, but rather an involved insistent *way* of responding and playing, so of finding things and dealing with them, on pain of 'giving up the game'." (342.)

[14] Haugeland claims that "skill repair is just the sort of multi-factor trading off or equilibrating that is familiar from holist epistemology"(336), but holist epistemology lacks the kind of inegalitarian structure Haugeland proposes. In Neuratian terms, Haugeland's theory seems to be that we are adrift on leaky ship at sea, a ship whose 'core' planks cannot be moved or replaced. If the remaining planks cannot be rearranged around this core in a fashion

On such an account, our constitutive skills regarding say, cats, could never turn out to incorporate a false set of standards. Our constitutive skills may tie mundane 'cat' findings to mundane 'animal' findings, and this combination might turn out to be unworkable if we consistently found all 'cats' to have always been robots.[15] According to Haugeland's view, however, this would only show that there were, in fact, no cats (we would have to "give up the game" of using 'cats' to understand the world).[16] However, even if we had to revise the constitutive skills associated with 'cat' when we started treating them as machines, in not turning out to animals, cats did not thereby 'annihilate themselves'.[17] Similarly, lilies did not annihilate themselves when they turned out not to be a natural kind, we just discovered that one of the standards we thought about lilies in terms of was incorrect.[18]

The dangers of overvaluing the constitutive rules shows up in Haugeland's general discussion of the "stances" associated with various sets of constitutive skills.

> a stance is more than just an attitude towards or a perspective on things, more even than a method and terminology for dealing with them. Adopting a stance is *taking a stand*. Why? Because it is this alone -- *commitment* to constitutive standards -- that allows that towards which the stand is taken to stand out as phenomena, to stand over against us as objects. Such standards determine the *being* of objects: what it is for them to be, *and* what is possible and impossible for them. Practitioners' insistence that the objects accord with the standards presupposes an ability to tell whether they do, and a resolve not to stand for it if they don't -- either by finding out what went wrong and fixing it, or (failing that) by "giving up the game". (284)

---

which will make the ship seaworthy, then we should simply 'jump ship' and board a new vessel with a different type of core.

15 A possibility discussed in Putnam 1975.

16 Katz 1975, draws much the same conclusion with respect to the possibility Putnam brings up.

17 The talk of objects annihilating themselves also suggests that they actually existed until the unworkability of the scheme was actually discovered; constituted objects exist "as long as [the equilibrium between the constitutive and mundane skills] lasts." I'm not really sure what to make of this, since it seems to suggest that there actually were (insert your favorite ultimately unconstitutable object here) until it was discovered that there couldn't be.

18 Much the same could be said of Haugeland's discussion of 'things'.

> "things" that seem to pop into and out of existence, to change their properties radically and capriciously, to differ in character depending upon who looked at them or how, or to be accessible only to specific individuals or under peculiar circumstances -- such "things" would be (ceteris paribus) repudiated as unreal. (350)

Admittedly, we would be inclined to think that some isolated phenomena was illusory if it seemed to pop in and out of existence, changed it properties, could be seen by some but not others etc. Nevertheless, if such a thing hung around and had some causal effects upon its surroundings, we might just think that some of our basic beliefs about things were, not that things had all 'annihilated themselves' and that there were, in fact, no things. Einstein might have thought that 'action at a distance' was part of the 'excluded zone' for physical objects, and this belief of his might be part of the common sense framework within which we understand objects. Still, a physicist can take himself to have discovered physical phenomena within this zone, and the phenomena are no less physical for it.

The stance that most interests Haugeland here is the 'intentional stance,' and how the norms of rationality can be taken to be constitutive of beliefs.

> The notions of 'commitment', 'allegiance', and 'proper evidence' are used here [by Davidson] without elaboration or defense; yet they suggest a deep (and, I think, deeply right) intuition to the effect that different schemes or realms entail different *standards* or criteria of adequacy to which descriptions of the phenomena must 'live up'. Or, to put it another way (accepting Davidson's connection between vocabulary and domain): entities themselves must live up to these standards, if they are to count as entities in this domain at all. Thus being rationally related to other mental states and events is a standard that any proposed candidate must meet if it is to join the mental club; being interrelated with others according to strict causal laws is the analogous entry condition for physical phenomena. (Haugeland 1992, p.29).

While it may be true that there is a close conceptual connection between belief and the standards of rationality, if rationality is understood as being constitutive of the mental in this fashion, one ends up with disastrous results.

If the norms of rationality are taken to be constitutive of beliefs, then, just as there cannot ('in the strict sense') be illegal chess moves, there cannot be sets of beliefs that are not in accord with the principles of rationality. Since we frequently criticize others for being irrational, this conclusion is highly paradoxical. Haugeland, however, having already accepted its chess analog (332), sounds, at times, as if he is willing to accept this paradoxical conclusion for beliefs as well.

> for intentional attribution, we *insist upon* rationality: in confronting apparent breaches of the standard, we first attempt to rectify them (e.g. explain them away in terms sanctioned by the standard itself); and, failing that, we give it up. (284)

This does not strike me as an especially attractive bullet to bite, and the case of belief highlights the vague dissatisfaction with Haugeland's attempt to neutralize the paradox in the case of chess. It is not enough to say that people with sets of beliefs that violate the constraints of rationality are *conceivable* but not (strictly) *possible*.[19] *Actual* belief sets almost invariably contain inconsistencies, and any account of the norms of rationality that denies this is in serious trouble.

Furthermore, this giving an 'analytic' status to the standards embodied in the constitutive skills leads to Haugeland's endorsement/explication of Kuhn's remark that "the proponents of

---

[19] "There must be two distinct senses of 'can' and 'possible' associated with any constituted domain . . . . The wider sense, which we might call the *conceivable*, comprises everything that the players, *qua* players, would have the resources to recognize or otherwise cope with, were it to occur. . . . The narrower or stricter sense, for which we can now reserve the term *possible*, includes only that which would accord with the constitutive standards, were it to occur. Thus, in this sharpened usage, illegal moves are conceivable, but not possible." (332.)

competing paradigms practice their trades in different worlds."[20] Haugeland claims that "the 'worlds' are different because the constituted objects are different; that is, they belong to different domains" (352). If one treats the constitutive standards as 'analytic', then, of course, the two-incommensurable-worlds view can seem like an unavoidable consequence of major changes in centrally held beliefs that have taken place over the history of science. However, if we really understand, say, both the Aristotelian and the Galiliean as trying to figure out what the empirical "game" is and how to play it, then we should allow that their constitutive skills really were fallible. Consequently, we could understand at least some of their mundane skills as shared (which they intuitively seem to be).

Consider, after all, Haugeland's discussion of 'disclosure':

> Finding that a game is playable – or, better, finding a game that is playable -- is therefore a kind of *achievement*, one that includes an element of "discovery" about the world. This element of discovery is quite different, of course, from the sort of discovery that is possible *within* a game: that one's rook is indefensible, or that the batter is trying to bunt, for instance. Rather, it's something like a prior or "meta" discovery to the effect that indefensible rooks or attempted bunts are ineligibly possible at all, and how they are possible. If we call such "meta discovery" *disclosure*, we can reserve 'discovery' for ordinary or mundane findings within the game. What is disclosed, then, is the playable game as such, including the intelligible domain within which its phenomena make the sense they make. (331.)

Treating people working with different paradigms as working with different objects might be acceptable if the two sets of constitutive skills really did 'disclose' two different sets of objects. However, the Aristotelian 'world' is not, ultimately, 'disclosable.' After all, that is precisely what, according to Kuhn, brings a paradigm to a 'crisis': the resilient finding that the world is not, ultimately, disclosable under that paradigm's set of constitutive skills. If the objects in the Aristotelian's world really were constituted by their constitutive skills, then the Aristotelian's weren't really practicing their trades in a 'world' at all. They walked around in *our* world but all *their* terms turn out to be about constructs in an elaborate, but undisclosable, fantasy. Giving constitutive standards the quasi-analytic status Haugeland wants for them makes this kind of radical detachment from any disclosable world quite likely. This seems to be, to say the least, a very unappealing consequence of his view.[21]

---

[20] Kuhn 1962, p 150.

[21] One can, of course, still allow a more pluralistic 'many worlds' view (in the Nelson Goodman's sense), since it may be the case that the world of quarks, the world of tennis rackets and the world of good and bad art are all,

Of course some people are perfectly satisfied without objectivity at the level of constitutive rules. Nevertheless, once we move from games to more empirical phenomena, the idea that both types of skills are up for correction seems fairly natural. Or, if it isn't 'natural,' it is, at least, more in keeping with the vaguely holistic climate that has characterized philosophical thought about our beliefs over the past 50 years.[22] As a result, if Haugeland wants to pass over more thoroughly holistic positions in favor of his neo-conventionalist alternative, he should give us some reason for doing so. All the concerns he uses to motivate his view (the need for objectivity etc.) seem to favor holist accounts (which could account for the objectivity of both sorts of skills) over conventionalist ones (which just accord objectivity to the mundane skills). Consequently, given his own reasons for introducing the account, an argument for stopping short of the holistic and fallibilist position seems called for.

ultimately, disclosable. But accepting this kind of many-worldism doesn't require accepting the type suggested in the Khun quotation.

22 Not that this climate has lacked its detractors (see, for instance, Fodor and Lepore 1992). But these detractors typically try to give an account of intentionality in terms of just the mundane skills in precisely the way Haugeland argues to be impossible.

References

Boghossian, P. 1989, "The Rule Following Considerations" *Mind*, **98**, pp.507-49.

Boghossian, P. 1997. "Analyticity" in Hale and Wright (eds.*) A Companion to the Philosophy of Language*. Malden: Blackwell, 1997.

Coffa, A. 1991 *The Semantic Tradition from Kant to Carnap*. New York: CUP.

Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge: MIT.

Fodor, J. 1990. *A Theory of Content*. Cambridge: MIT.

Fodor, J. and LePore, E. 1992, *Holism: a shoppers guide* Cambridge: Blackwell.

Haugeland, John, 1992 "Dasein's Disclosedness" in Dreyfus, H. & Hall, H. (eds). *Heidegger: A Critical Reader*. Cambridge: Blackwell 1992.

Haugeland, J. 1998. *Having Thought*. Cambridge: Harvard University Press.

Katz, J. 1975. "Logic and Language: An Examination of Recent Criticisms of Intentionalism", in Gunderson, K. (ed.). *Minnesota Studies in the Philosophy of Science VII*, Minneapolis: University of Minnesota Press, 1975.

Kripke, S. 1982. *Wittgenstein on Rules and Private Language*. Cambridge: Harvard University Press.

Kuhn, T, 1962. *The Structure of Scientific Revolutions*, Chicago, University of Chicago Press.

Loewer, B. 1997. "A guide to naturalizing semantics" in Hale, B. and Wright, C. (eds.*) A Companion to the Philosophy of Language*, Cambridge: Blackwell, 1997.

Milikan, R.G. 1984. *Language, Thought, and Other Biological Categories*, Cambridge: MIT.

Putnam, H. 1975. "The Meaning of 'Meaning.'" reprinted in his *Mind Language and Reality*. New York: CUP, 1975.