The Informational Turn in Philosophy[1]

FREDERICK ADAMS
Department of Cognitive Science & Philosophy, University of Delaware, Newark, DE 19716, USA; E-mail: fa@udel.edu

Abstract. This paper traces the application of information theory to philosophical problems of mind and meaning from the earliest days of the creation of the mathematical theory of communication. The use of information theory to understand purposive behavior, learning, pattern recognition, and more marked the beginning of the naturalization of mind and meaning. From the inception of information theory, Wiener, Turing, and others began trying to show how to make a mind from informational and computational materials. Over the last 50 years, many philosophers saw different aspects of the naturalization of the mind, though few saw at once all of the pieces of the puzzle that we now know. Starting with Norbert Wiener himself, philosophers and information theorists used concepts from information theory to understand cognition. This paper provides a window on the historical sequence of contributions made to the overall project of naturalizing the mind by philosophers from Shannon, Wiener, and MacKay, to Dennett, Sayre, Dretske, Fodor, and Perry, among others.

> At some time between 1928 and 1948, American engineers and mathematicians began to talk about 'Theory of Information' and 'Information Theory,' understanding by these terms approximately and vaguely a theory for which Hartley's 'amount of information' is a basic concept. I have been unable to find out when and by whom these names were first used. Hartley himself does not use them nor does he employ the term 'Theory of Transmission of Information,' from which the two other shorter terms presumably were derived. It seems that Norbert Wiener and Claude Shannon were using them in the Mid-Forties.
>
> (Yehoshua Bar-Hillel, 1955)

1. Introduction

> ...the amount of information is defined, in the simplest cases, to be measured by the logarithm of the number of available choices. It being convenient to use logarithms to the base 2... This unit of information is called a 'bit' ... a condensation of 'binary digit'. (Warren Weaver, 1949)

The information age, in the technical sense of information derived from the mathematical theory of communication (Shannon and Weaver, 1949), began around the year 1950 (plus or minus two).[2] There is not much doubt that the information age and the computing age developed in concert, and that is why I set the beginning at 1950 (the year of Turing's famous paper on computing and intelligence). There were, of course, forerunners[3] of the development of information theory in several disciplines and I cannot hope to retrace all of the steps here[4]. So for the purposes of this survey of the application of information theory to problems of philosophy, I will simply declare the midpoint of the Twentieth Century as the beginning of the informational turn.

Minds and Machines **13:** 471–501, 2003.

It is also the case that information theory has touched many areas of philosophy including, computation theory, artificial intelligence research, perception, knowledge, logic, philosophy of mind, philosophy of language, philosophy of science, decision theory, and philosophy of music – among others. Unfortunately, due to limitation of space, I must choose but one strand to follow here. Since the one I know best is the one that marches through the philosophy of mind and language, that is the one I shall trace. Even then what follows will barely scratch the surface[5] of what has been done in the application of information-theoretic concepts to these philosophical issues since 1950. I will limit my discussion to primary sources that attempt to solve philosophical problems through the application of concepts from information theory proper. And I shall focus mainly on the build-up to a naturalization of the mind and meaning.

At the foundation of information theory is the development of methods to measure the amount of information generated by an event or events, and mathematical treatments of the transmission characteristics of communication channels. Shannon (1949) beautifully laid out and described mathematical concepts of informational *measures, sources, receivers, transmitters,* concepts of *noise, uncertainty, equivocation, channel, channel capacity,* and the characteristics of *continuous* vs. *discrete information channels*. I shall briefly touch upon some of the details below. As Shannon saw it "[t]he fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point" (Shannon and Weaver, 1949, p. 31). It is easy to see why someone like Dretske (1981, 1983) would be interested in applying information theory to knowledge,[6] when the source of Shannon's "message" is the world and the receiver is the mind of a would-be knower. If we could "reproduce exactly or approximately" the message the world sends we could acquire knowledge of the world. Informational concepts may also help solve long-standing problems of causal deviance, causal over-determination, and other problems associated with causal theories and counterfactual theories of perception and knowledge.[7]

Though the application of information to knowledge may have been easier to see, it took longer to recognize that information theory might be useful in tackling problems of minds and meaning. This was due in no small part to Shannon's claim: "Frequently the messages have *meaning*; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem" (Shannon and Weaver, 1949, p. 31). For the reason Shannon mentioned, perhaps philosophers interested in what would become the *cognitive sciences* were slow to see through the engineering problem to the usefulness of the formal features of communication theory.[8] Eventually they would see that information theory could supply tools needed to help solve philosophical problems of perception, cognition, and action. If one thinks of the mind as receiving information from the environment, storing and coding that information, and then causally guiding behavior in virtue of the stored representational content, it is not too hard to see why information theory would apply to these elements. Matters of causal deviance may be resolved by requiring unequivocal, noiseless communication channels for knowledge and action. Matters of meaning may be resolved by showing how semantic content of cognitive states derives from informational origins. I believe that the move from information to meaning, using informational concepts, is one of the major successes of the last half of the 20th Century in philosophy. So I shall trace steps along the way to that end in this essay. Several different people saw parts of the overall picture and contributed significantly to the events to unfold. I shall discuss several of the key figures contributing to the informational turn and describe some of their important contributions. For the most part, I will do this in an historical order, beginning with the inventors of the concepts.

Wiener (1943, 1948, Chapter 4) himself applied information theory (cybernetics) to the explanation of purposive controlled behavior in feedback systems, by looking at such systems gone wrong in cases of *ataxia* and *tabes dorsalis*. In the latter, syphilis destroys the ordinary information feedback through the spinal cord thereby interrupting information conveyed via kinesthetic sensations. Joint and tendon signals are not properly processed and information concerning posture and smooth motion is interrupted — patients loose proprioceptive or kinaesthetic sense. Wiener was keenly aware that such maladies were symptomatic of failed information feedback control systems in the brain. His mathematical theory of *cybernetics* (a term he supplied) attempted to capture the formal structure of the informational breakdown in such systems. However, it is important to note that his is a treatment of the *amounts* of information and the mathematics of the time series and oscillations that may be applied to the nervous system, and not of the *contents* of such information signals being processed. What I believe was the first plausible account of that would come thirty years later.

When there is no damage and everything goes smoothly in the control systems of the body, Wiener notes that "[w]e do not will motions of certain muscles, and indeed we generally do not know which muscles are to be moved to accomplish a given task...say, to pick up a cigarette. Our motion is regulated by some measure of the amount by which it has not yet been accomplished....The information fed back to the control center tends to oppose the departure of the controlled from the controlling quantity, but it may depend in widely different ways on this departure" (Wiener, 1948, p. 116). Of course, when things go wrong, control systems fail and bodily movements oscillate out of control. In any case, Wiener and colleagues, as early as the first days of the information-theoretic age, had begun to apply informational concepts to important concepts in the philosophy of mind — to the analysis of teleological behavior of purposive systems.[9]

Also, like Turing (1950), Wiener thought that the brain was an information-processing computer with implementation of logical inferences, memory storage, and so on. Wiener was keenly aware of the logical, computational, and biological properties necessary to begin to model thought as computation in the brain. However, so far as I can tell, he had no idea of how the contents of propositions would be acquired or stored. That is, he had not yet embarked on a project to fully naturalize minds or meaning.

Wiener spoke briefly and indirectly of the topic of pattern recognition and an attempt to provide a cybernetic model (1948, Chapter VI) of "Gestalts and Universals." He noted that there are feedback systems that focus the eyes on targets, causing them to foveate. He also believed retinal and cortical firing patterns could be modeled along cybernetic lines to explain the "gestalt" phenomena of pattern recognition, but his description was only a beginning of the work to be done by people such as Gibson (1979) and Marr (1982). He inspired a beginning of the information processing approach to vision and pattern recognition (Sayre, 1965, 1976). He saw a role for cybernetics in as diverse areas as psychopathology and the explanation of meaning in language and mind, but merely made the initial investment upon which dividends would be paid later in the century.

2. Information Is Not Meaning

> The trouble here appears to be due largely to a confusion of the concept of *information* with that of *information-content* — the confusion of a *thing* with a *measure* of a thing. Communication engineers have not developed a concept of information at all. They have developed a theory dealing explicitly with only one particular feature or aspect of messages...their unexpectedness or surprise value. (MacKay, 1969, p. 56)

Bar-Hillel (1955) was perhaps the first philosopher to stress that the word "information" exploits an ambiguity between frequencies, frequency dependencies of events, and meaning. Hartley (1928) was a precursor to Shannon. He may have been the first to introduce the logarithmic measure of *amounts* of information, the same measure with which Shannon begins. Bar-Hillel points out that Hartley's "new insight...was that the measure he was looking for had to be dependent upon the 'frequency relations involved in electrical communications' " (1955, p. 93). He impresses upon us that "transmit information" as used by Hartley, "has certainly nothing to do with what we might call the *semantic content* of the signals transmitted" (1955, p. 93). About Hartley's work, Bar-Hillel continues: "when he speaks of the 'measure of information,' or the 'amount of information,' or 'information content' in a signal sequence, he has nothing else in mind than a certain function of the relative frequency of this sequence among the set of all possible signal sequences of the same length" (1955, p. 93).[10] Indeed, there was a fairly wide-spread disdain for the need to associate what philosophers today think of as semantic content with the transmission of signals and carrying capacities of communication channels, in the early days of information theory. Bar-Hillel desired to change that.

Bar-Hillel (1955) points out that "the event of transmission of a certain statement and the event expressed by this statement are, in general, entirely different events, and the logical probabilities assigned to these events, relative to certain evidences, will be as different as will be the frequencies with which events of these kinds will occur relative to certain reference classes" (1955, p. 96). He concludes: "...the concept of semantic information has intrinsically nothing to do with communication" (1955, p. 96).

Bar-Hillel complained that Wiener occasionally moved from "information" to "meaning," without comment. He contrasts this with Shannon's insistence that the semantic aspects of communications are "irrelevant to the engineering problem." Bar-Hillel himself shared the optimism of Warren Weaver who believed that Shannon's work was relevant to the semantic aspect of information (Bar-Hillel, 1955, p. 98), but who did not say how.

Bar-Hillel's and Carnap's own semantic theory of information was not naturalistic. It presupposed the existence minds and language, and did not use the concept of information to explain how these came to be: "The theory we are going to develop will presuppose a certain language system and the basic concepts of this theory will be applied to sentences of that system. These concepts, then, will be semantic concepts, closely connected with certain concepts of inductive logic..." (Bar-Hillel, 1964, p. 221).

They also proclaimed that their account would not deal with what Warren Weaver called the "semantic problem of communication" and which is concerned "with the identity, or satisfactorily close approximation, in the interpretation of meaning by the receiver, as compared with the intended meaning of the sender" (1964, p. 222).

They did however develop a semantic theory in so far as it embraced the truth-values of sentences and their logical relations. Two of their most controversial claims are that mathematical and logical truths generated *zero* information (in the sense of the amount of information generated by things that could not be false) and that self-contradictions assert *too much information to be true*. The consequences of these claims are still being felt and debated (Floridi, forthcoming a, Fetzer, forthcoming).

Since I am detailing here the steps that lead to a use of informational concepts to naturalize meaning and purposive activity, I shall turn to some of the basic features of information that would go into such accounts.

To detail how information is relevant to meaning requires seeing how to isolate a specific piece of information carried by a signal and to feature that piece of information as the semantic content of a cognitive structure. That piece of information has to generate a meaning, something that can be falsely tokened. As nearly as I can tell it was Fred Dretske (1981) who first fully saw how to connect the dots while utilizing the concepts of information theory.[11] Others saw steps along the way. For instance, Paul Grice (1957) saw the link from natural meaning to what he would come to call *non-natural meaning* (and what I am calling here *semantic content*). Grice's *natural meaning* was a natural sign or indicator (smoke in the forest naturally means or indicates the presence of fire). Non-natural meaning is what I'm calling semantic content (the word "smoke" does not naturally mean or indicate fire, but it does semantically mean *smoke*). I will complete the story of the transition from information to meaning, when I get to Dretske's account of this transformation (below).[12]

Here I begin with some preliminaries about information as a commodity dealing in probabilities of events and relationships between those events. To be of value to a would-be knower, or to someone interested in naturalizing the mind, information must be an objective, mind-independent commodity. In principle, it should be possible for someone to be the first person to learn that p. If S were the first person brought to know that p by the information that p, then the information that p would appear to have objective properties. The following examples suggest that this is so. Waves of radiation traveling through space may contain information about the Big Bang before anyone detects it. Fingerprints on the gun may contain information about who pulled the trigger before anyone lifts the prints. Thus, information appears to be mind-independent (and, thereby, language independent too).

Information must also be capable of having a very special relationship to the truth. Since one cannot know what is false, if information is going to bring one to know that p, then information must also be tied to the state of affairs that makes p true. Otherwise, it is hard to see the value of information over belief and truth itself. On at least some accounts, information has this connection to truth (Bar-Hillel, 1964; Dretske, 1981; Floridi, forthcoming a, b). One can be misinformed. One can be informed that q, when one needs to know that p, but one cannot be misinformed that p. For something can only carry the information that p, if p. Indeed, if we think of information as being contained or carried in one event (set of events) and as being about another event (set of events), then the transmission of information is the product of a correlation and dependency between the two events (sets). To see this in more detail, let's consider Dretske's (1981) attempt to explicate the Shannon-Weaver account of information.

To adapt information theory to a format friendly to a theory of mind or knowledge, several matters need to be resolved. For example, to know that Bush was elected president involves information being generated by the event of his election. It also involves transmission of that information to a prospective knower S. S must detect physical events that carry that transmitted information, and those events must cause or sustain S's belief that Bush was elected.

Let's begin with generation of information. An event's occurrence generates information. How much is generated is a function of how likely was the event's occurrence. The more likely an event, the less information it generates — while the less likely the event, the more information it generates. For example, on any random day, telling you truly that it is going to rain today is more informative in Phoenix than Seattle. Different ways of classifying events may result in different amounts of information generated. And there are many different ways of trying to measure or quantify amounts of information. Dretske follows the communication industry standard (Weaver and Shannon, 1949) of measuring information in bits (*bi*nary digi*ts*), representing the number of binary partitions necessary to reduce a collection of equally probable outcomes to one (e.g., beginning with 8, a three-step reduction to 4, to 2, to 1 = 3 bits). The amount of information generated at a source s by the reduction of n equally likely possibilities to one is represented: I(s)=log n (base 2). Here I(s) represents the average amount of information generated at a source by a reduction of equally likely events. If the range of possible events at the source $s_1, s_2, \ldots\ s_n$, are not all equally likely, then the amount of information generated by the occurrence of s_I is: $I(s_I) = \log 1/p(s_I)$ (where p = probability). So, for example, suppose ten persons apply for a job and nine are from inside the company, one from outside. If s_1 is the selection for the job of someone outside the company, then $I(s_1) = \log 1/0.1 = 3.33$ bits of information. For contrast, selection of someone from inside the company, s_2 would generate 1/0.9 = 0.15 bits of information.

Next, let's consider information flow or transmission. For information at a receiving point r to be about a sending point s, there must be dependence between the events at r upon those at s. Suppose at s there are eight candidates equally likely to be selected. A selection of Susan generates 3 bits of information. Suppose at r there are eight equally likely names that may be put on the employment forms in the employment office. A selection of "Susan" generates 3 bits of information. But there would also be 3 bits generated if, mistakenly, the name "Tony" were placed on the employment forms. Clearly, though this amount of information is the same, it is not the information that Susan was selected. We want the information at r to be about the events that transpired at s. Letting "$I_s(r)$" represent this information, $I_s(r) = I(r) -$ noise. Noise is the amount of information generated at r that is independent of what happens at s (not about s), and when "Tony" is placed on the forms, but Susan was selected, the noise = 3 bits. Thus, no information about s arrives at r.

Now for our purposes, the import of these formulae for calculating amounts of information is not so much the absolute values of information generated or transmitted by an event, but the conditions necessary for transmission. For most events it would be difficult or impossible to determine the exact probabilities and ranges of possibilities closed off by an event's occurrence. What is important is whether one receives at r as much information as is necessary to know what happened at s (under a relevant specification). For a signal or message to carry the information that Bush was elected, it must carry as much information as was generated by Bush's election. We know this is more information than that a Republican ran for office, and more than that someone was elected. Calculating exactly how much information is generated by Bush's election is not as important as determining under what conditions the information that does arrive carries the information that Bush was elected. This is what Dretske calls the informational content of a signal.

Informational content: A signal r carries the information that s is F = The conditional probability of s's being F, given r (and k), is 1 (but, given k alone, less than 1). K is a variable that takes into account how what one already knows may influence the informational value of a signal. If one knew nothing, k would go to zero. If I know that Vice President Cheney is from Texas or Wyoming, and I learn that he is not from Texas, I thereby have the information that he is from Wyoming. If you hear that he is not from Texas, but don't already know Wyoming is the only other possibility, you do not thereby receive the information that he is from Wyoming.

This account of the informational content of a signal has important virtues. If a signal carries the information that Bush was elected, then since the conditional probability that Bush was elected, given the signal is 1, then *Bush was elected*. Hence, the account gives information a connection to truth. Clearly it will also be the case that the signal carries as much information about s, $I_s(r)$, as was generated by the fact that Bush was elected. Noise about the fact *that Bush was elected* is zero. Hence, the account gives us a way to understand transmission or flow of information of a specific propositional (factual) content from source to receiver—not just amounts of information.

However, we can now see clearly why information is not *semantic content* (or meaning). We can see this for at least two reasons. First, any signal that carries the information that *s* is *F* will carry the information that *s* is *F or G*. This follows form the fact that if the conditional probability of *s*'s being *F*, given some signal *r*, is 1, then *s*'s being *F or G*, given *r*, is 1 as well. So anything nomically, logically, or analytically entailed by being *F* will be nested in any information carried in a signal about something's being *F*. Hence, there is far *too much information* in any signal for the signal to come to mean only *F*, due to the information carried by the signal alone. Second, smoke (the stuff in the forest) in the right environment, carries the information that there is fire in the forest. But "smoke!" yelled out by me (or my thought that there is smoke here) need not carry the information that there is fire here. I may lie or be mistaken. Statements and thoughts can be falsely tokened. Still "smoke" semantically means smoke, falsely tokened or not. And my thought that there is *smoke* semantically means smoke, falsely tokened or not. How is that possible? How can a symbol come to have a semantic meaning, *smoke,* but not necessarily carry information about smoke? This is especially puzzling if the tokening of a symbol owes its *semantic content* (meaning) to its informational origins.

In Section 4, I will try to describe the jump from information to meaning. In Section 3 I want to say more about philosophers who saw the relevance of information theory to issues in the philosophy of mind. I do not think that they had the same picture of the transition from information to meaning as Dretske, but I do think that they all clearly saw that informational concepts should take center stage in the push toward understanding what a mind is and how the mind works.

3. Purposive Behavior and Cognition

> I have mentioned before that the only events in the universe which require the transmission of information are goal-directed activities.
>
> (Sommerhoff, 1974, p. 91)

D.M. MacKay in a series of papers in the 1950s applied concepts of information processing to goal-directed behavior. As was the case for Rosenblueth, Wiener, and Bigelow (1943), MacKay (1951, 1956) saw goal-directed systems as having an input to an organism or machine that represents the current state of the environment y and an input x that represents the environment plus the goal state A of the organism. Then there must be an effector that "gives rise to activity leading to the minimisation of some measure of xy" (MacKay, 1951, p. 226). Detecting the state of the environment and the output of the effector in the organism required information to flow into the organism. Unfortunately, as for Rosenblueth et al. (1943) requiring both x and y to be inputs from the environment guaranteed that the model could not be applied to systems that seek goal-states that do not now (or maybe ever) exist, but the model was good at modeling munitions homing in on a target. And any improvement on the model would await a naturalized theory of content and representation.

MacKay was well aware of the problems of stability in "error-operated feedback systems," and importantly, he was aware that the *information* in such feedback systems was the same commodity as was being discussed in the fledgling theory of information taking shape around him. Perhaps most importantly, for MacKay, is that his sense of information becomes tied to an operational notion of information because he was interested in the information's use to the goal-directed system receiving it. MacKay (not unlike Dennett) sees information in terms of its potential to impact the "conditional readiness" of an organism or system. "It is the hierarchy of such readinesses — my total state of readiness for adaptive goal-directed activity — which changes when I gain information. Information in fact could be defined in actor-language as that which alters my total adaptive readiness in this sense" (1969, p. 60). Now one might think that knowledge is power and that MacKay is just pointing out that when one gains information about the world that information could be used to enhance one's adaptive readiness. He considers that view, but seems to reject it. His is a view that maps outputs of the body and changes in the environment onto perceptual changes of the world's events. "The basic symbols in which our 'model' of the world could most economically be described would stand, not for objects in the world, but for characteristic patterns in the events of perception. These events themselves, be it remembered, we have taken to be acts of internal adaptive response" (1969, p. 61). For MacKay the relationship between "information" and adaptive response appears more intimate than one might otherwise suppose: " 'Information about-X' is that which *determines the form* of one's state-of-readiness for X" (1969, p. 107). (In keeping with my earlier claims, I don't find anywhere in MacKay's work a good notion of what makes something "about X" in a semantic sense of "about.")

Curiously, MacKay uses 'information' and 'meaning' interchangeably. For note that he says virtually the same thing about meaning that he says about information. "It looks as if the meaning of a message can be defined very simply as the selective function on the range of the recipient's states of conditional readiness for goal-directed activity: so that the meaning of a message to you is its selective function on the range of your conditional readiness. Defined in this way, meaning is clearly a relationship between message and recipient rather than a unique property of the message alone" (1969, p. 24). It is pretty clear that he does not mean by "mean" the *semantic content* of a signal in the sense of something that could be *falsely tokened*.[13] Equally revealing is the following pasage:

> I have suggested elsewhere that the internal representation of the world by an organism may be thought of as a statistical model of the 'pattern of demand' made by the world on the organism. By the 'pattern of demand' I mean not merely those features of the world (such as heat and cold) that bear upon and disturb equilibrium of the inert organism, but all of those that the active organism has to take into account when conducting goal-directed activity. The suggestion is that the organising system developed to match this pattern of demand (to do the necessary 'taking into account') can *itself* serve as the internal representation of the world. (MacKay, 1969, p. 112)

MacKay saw learning in such a system as a "process...in which the frequency of past successes or failure of a given action determined the transition-probability to that action in the future" (MacKay, 1951, p. 231). He also discussed models of pattern recognition that were similar to those put forth by Wiener and much later by Sayre (1965, 1976). These models were what he called "template-fitting" models (see Sayre below and Sayre, 1986). MacKay (1969) was also taken with the threat that mechanism of the brain might rob the brain of meaning. In the end, wisely MacKay would claim this was a false dichotomy (mechanism vs. meaning).

I should also point out that MacKay (1969) was also interested in an information based account of what makes something a statement, command, question, and so on. He was likely the first person in information circles to work on what would become "pragmatics."

Wooldridge (1963) was impressed with the similarity between the natural and adaptive processes that support the cognitive and mechanical processes of the brain and the operational principles of man-made devices. In particular, he was impressed with the similarity between the brain and a digital computer, saying "an essential property of the whole nervous system is that it transmits information by electrical means and that the type of electrical conduction ...is of an all-or-nothing nature. It is as though the basic mechanism of the nervous conduction consisted of some form of electronic on/off switch!" (1963, p. 5). Wooldridge, as were many others, was especially keen to compare the synaptic firings of the brain to logic circuits, such as *and*,*or*, and *not* gates.[14] Like Dennett (1969), he was impressed with the functional decomposition of complex cognitive tasks into simpler and simpler tasks until finally the sub-functions are realized by teams of on/off switches (1963, p. 234). However, Wooldridge did not appear to see how to utilize the technical notions of information to analyze the matters of content, representation, or knowledge.

Armstrong (1968, p. 139) appealed to the general notion of information and negative feedback in characterizing purposive activity, but curiously he uses the notion in a way that later users would not. He thinks information helps solve the problem of the intentionality of behavior. S may do A under description "A" intentionally, but not under description "B" even though A = B. Armstrong implies that S may not be receiving the information that S is doing B. But, of course, if one is receiving the information that A is happening, then one *is receiving* the information that B is happening, when A = B. It is true that *the way* a signal carries one piece of information may not tip a user that it also carries the other piece of information (even though it does). It was surely this feature that Armstrong was after. Armstrong adds that S may not be perceptually aware that B is happening, and this should be the heart of the matter. It is a difference of perceptual recognition, not of information, that would be able to account for the difference.

Armstrong sees clearly that there is an important relationship between information and conceptual deployment, in this work and in this example. This association clearly would continue in Armstrong's later work. He even has a notion of something like a concept's "locking" to a piece of information. He claims that *a structure's carrying information* has to be able to have causal powers – saying that the information in a cause must be able to "turn off the power" (or turn it on) in a mental cause (such as a purpose).

Armstrong's primary discussion of information is in regard to perception. He says "A perception which involves an inclination, but no more than an inclination, to believe, may be conceived of as the acquiring of information which we have some tendency, but no more than some tendency, to accept" (1968, p. 225). Here he is working the difference between perception (which he thinks involves belief) and something less (what he calls an "idle perception"). He calls perception "a flow of information" (1968, p. 226), and "perceptual experience" as opposed to "mere perception" is "this flow of information in so far as we are conscious of it...introspectively aware of it" (1968, p. 226).

Since Armstrong sees clearly that information is relevant to the analysis of goal-directed behavior (1968, p. 255), it is very likely that he knew of the work of Wiener or Shannon, but they are not referenced in his 1968 book where he specifically discusses a role for information.

Dennett (1969, 1981, 1987) was among the first philosophers to give informational concepts more than a cursory nod. In the development of what we now know as Dennett's *intentional stance* theory of the mind, informational concepts played a prominent role. From the start Dennett (1969, pp. 45–47) maintained that "[n]o creature could exhibit intentional behavior unless it had the capacity to store information." He went on to characterize a computer's intelligent store of information as illustrated by the capacity to produce a sequence of characters in response to a particular cue. "Indeed this storage can be called information storage only by grace of the fact that the users of the output can interpret it as information....We should reserve the term 'intelligent storage' for storage of information that is *for* the system itself, and not merely *for* the system's users or creators. The criterion for intelligent storage is then the appropriateness of the resultant behaviour to the system's needs given the stimulus conditions of the initial input and the environment in which the behaviour occurs."

Dennett realized clearly that information was not captured by the physical properties of inputs alone. If the brain cannot react differentially to stimuli in appropriate response to environmental conditions they herald, it will not serve the organism at all. But how is the brain to differentially respond to inputs on the basis of more than their physical properties alone? "No physical motions or events have intrinsic significance...the capacity of the brain to discriminate by significance cannot be simply a capacity for the analysis of internal structure, electro-chemical or cryptological, of the input sequences. (1969, p. 47)"

Dennett (1969, p. 55) was also aware that "[n]o sense has yet been given to the claim that a neuron's impulses are signals with content or meaning, but if, for example, a particular neuron in the optic nerve fires its output if and only if there is a particular pattern of stimulation on the retina (due to the particular summing effects of the neurons in the lower ranks leading to its input), in a borrowed sense one could say that the neuron's output is unambiguous." Dennett explains how large scale mechanisms of redundancy in the brain and selectional mechanisms at the level of the species (inter-cerebral) and the individual (intra-cerebral) can help to reduce the types of ambiguities we know to exist. Hence, Dennett was trying to understand the types of environmental constraints that must be in place to have non-equivocal transmission of information from the environment to the cognitive agent and within the brain of the agent. And he was trying to understand *how* a system was able to mine the *informational* benefit of the inputs when confronted only with the physical properties of the inputs.

At even this early stage of his work, Dennett shies away from taking *information* to be a real, mind-independent commodity. He introduces a doctrine I will call (using his words) "discrimination by significance." Dennett puts it this way: "...since a stimulus, as a physical event, can have no intrinsic significance but only what accrues to it in virtue of the brain's discrimination, the problem-ridden picture of a stimulus being *recognized* by an animal, meaning something *to* the animal, prior to the animal's determing what to do about the stimulus, is a conceptual mistake (1969, p. 76)." Dennett continues: "The criterion for intelligent information processing must involve this behavioral link — however mediated — since propitiousness or adaptiveness of behavior is at least a necessary condition of intelligence" (1969, p. 77). In these early remarks by Dennett, one cannot help but see the "operational" view of information of that we just saw in the words of MacKay.[15]

In this early work, Dennett also revealed his trademark "heuristic overlay" view of content: "The ideal picture, then, is of content being ascribed to structures, events and states in the brain on the basis of a determination of origins in stimulation and eventual appropriate behavioral effects, such ascriptions being essentially heuristic overlay on the extensional theory rather than intervening variables of the theory" (1969, p. 80).

One might think that he is here talking only about *semantic content* of the type appropriate to ascribing propositional attitudes from the intentional stance. But the heuristic overlay appears to go deeper — all the way down to the level of *information* itself. "Information is not preserved in a sentence like a fossil in a rock; a sentence is a vehicle for information only in that it is part of a system that necessarily includes sub-systems that process, store, and transmit information non-linguistically. (1969, p. 88)" "...something is a message or a signal only when it goes on to effect functions in some self-contained intentional system. (1969, p. 186)." Clearly, by this account, one could not use information to explain the origin of the intentional mind...for it would be circular. And further, one is never given an account of the origin of the content of the heuristic overlay. That is, one never receives from Dennett an account of the origin of the content *within the intentional stance*. These features of his work have held constant pretty much throughout his career, and we can see them clearly in this, his first major work.

From the start, Dennett knew some of the details[16] of the mathematical theory of information and demonstrated that knowledge (1969, pp. 186–187). He knew that information transmitted was relative to knowledge of the receiver. He knew that information transmitted was relative to the possibilities determined by the description (partitioning) of the sending station and receiving station, and so on. Apparently, he believed this disqualified information from playing any constructive role in a naturalization project (though it is far from clear that it does). I (buying into the naturalization project) cannot help but find it ironic that Dennett was among the first philosophers to see the value of information theory for building the "intentional stance," while he himself never fully accepted that information theory could naturalize semantic content. Apparently Dennett has always believed that information is disqualified from the project of naturalizing the mind because informational content is expressed in propositional form. It seems a metaphysical aversion to *propositions* has been the stumbling block, all these years. "The information characterized by formal information theory, measured in bits or bytes...is hardly the concept we must appeal to when we speak of information-processing models of the nervous system. The information measured in bits is content-neutral....but despite ingenious (if unsuccessful) attempts to develop the requisite concept [of *semantic content*] as an extension of the information theoretic concept (Dretske, 1981; Sayre, 1986; Dennett, 1986), we still have no better way of individuating portions of the wonderful stuff than by talking in one way or another about propositions" (1987, p. 206). He concludes: "If the philosophical theory of propositions is as bad off as I claim, are these enterprises all in jeopardy..?" The implied answer is clearly *yes*!

Sayre (1965, 1969, 1976, 1986, 1987) clearly did want to naturalize the mind via information theory, contrary to Dennett. Sayre first applied concepts of information to recognition, then to consciousness and finally to a very wide range of philosophical problems indeed. Since I believe Sayre's last book on informational topics (1976) is his most developed and mature, I will focus mainly upon it.[17] However, I would be remiss not to point out that an earlier book (Sayre, 1969) has excellent discussion of the shortcomings of the paper by Rosenblueth et al. (1943) in accounting of purposive behavior.

Sayre (1976) expressly acknowledges profiting from the works of Armstrong (1968) and Dennett (1969). However, he explicitly chides them for "laxity in use of the term 'information' " (Sayre, 1976, p. xi). Sayre claims that Armstrong equates information and belief (Armstrong, 1968, p. 210), and treats knowledge as information about our environment, maintains that information is acquired by "bringing objects, events, etc., under concepts," and claims that information can be true or false (Sayre 1976, pp. xi–xii)." Sayre (1976, p. xii) applauds Dennett's introduction of information as diminished uncertainty, but also notes that Dennett thinks information can be "relevant," (Dennett, 1969, p. 170) to experience "in general" (Dennett, 1969, p. 150) or capable of being "true or false" (Dennett, 1969, p.157). As Sayre (1976, p. xii) correctly notes, "[n]one of these further senses is provided by communication theory, and most have semantic overtones which expositors of this theory often explicitly disavow." For his own account Sayre claims the virtue that "[i]n the text below, the concept of information is explicated in formal mathematics before being deployed in the analysis of other concepts, and the term 'information' is used only in senses that have been explicitly defined. Since information can be defined in this fashion, although fundamental it is not a primitive concept. The primitive concept in this treatment is that of probability, needed to interpret the formal definition of information. Although various analyses of probability are available from other contexts, none lends itself to this interpretation more readily than others. To accept probability as a primitive concept is to decline further attempts at clarification" (Sayre, 1976, p. xii). Sayre calls his view "informational realism," since he explicitly endorses information as an ontologically basic entity (Sayre, 1976, p. xiii).

Sayre wanted to use *information* as an ontological category to unite the categories of the mental and the physical, not unlike Russell's (1921) attempt to derive the mental and the physical from a "neutral" basis. Sayre says: "If the project of this book is successful, it will have shown not only that the concept of information provides a primitive for the analysis of both the physical and the mental, but also that states of information ...existed previously to states of mind." (Sayre, 1976, p. 16) Hence, Sayre clearly had the idea of information as an objective, mind-independent commodity. Indeed, Sayre believed Russell's approach to find a neutral basis from which to construct both body and mind failed because there was no "independent theory of sensibilia." Sayre believed his account had a better chance of succeeding because "...information, by contrast, is part of a firmly established formal theory, with proven applicability...." (Sayre, 1976, p. 17). He further believed that *information* would be more respected than Russell's *sense-data* because it came from a branch of mathematics, and was thereby better understood than even physics or chemistry. He also believed that since information would be ontologically basic, it would enable the explanation of how the mental and the physical interact causally (notoriously lacking in other theories or attempts at reductions). Most importantly, Sayre believed other attempted reductions moved from the *better understood* (the physical) to the poorer understood *sensibilia* (Russell) or *ideas* (Berkeley). This would not be the case with information as the ontological base. Sayre even went so far as to suggest the idea of explicating the concept of the physical in terms of information.

Sayre tries to model a long list of things as processes of information feedback, including: causation, life, evolution, evolutionary success, learning, and consciousness. I will focus only on the latter two. Sayre introduces a cybernetic model of learning and argues that it is broader than classical or operant models of learning and that the latter two are derivable from it. The model basically relates the probabilities of afferent states to effector states, showing how reinforcers and punishers can affect the probability dependencies between the two (Sayre, 1976, pp. 125–127). The model is based upon two general principles:

(I) Efferent states followed proximately by a reinforcing afferent state increase the probability of further occurrence up to a point of steady high probability, and that this probability subsequently decreases if the association between the two states discontinues.

(II) Efferent states followed by a punisher afferent decrease in probability of occurrence, more rapidly as their association is repeated.

The first postulate is designed to explain the acquisition of learned paired response and the second the possibility of its extinction. In combination, Sayre thinks the two can explain much complex learning (both classical and operant learning).

Perhaps what is striking about the model is that few formal features of information theory are employed — though it is clear that probabilities and their dependencies are involved in learning and extinction of behavior or patterns of neural firings. Also there is no appeal to the content of a piece of information in the learning model. There is learning (behavioral learning) but no learning *that p* in such a model (or at least not on the surface).

As for perceptual consciousness, Sayre conceives of it as "(very short-term) learning of the sensory system" and an "adaptation of the organism's afferent information-processing system...to environmental contingencies" (1976, p. 139). Sayre portrays the visual system as a "cascade of information channels, extending from the retina to the visual cortex and articulated at several junctures in between" (1976, p. 150). He treats perceptual fixation upon an object, such as a penny, as a process of both negative and positive feedback setting up an informational channel between retinal level B and visual cortex C. "The channel B–C is characterized by a high degree of mutual information as long as the regular configurations at level C accurately mirror the more dominant among the changing contours at level B" (1976, p. 151). We can see that Sayre is using cybernetics to model the mathematical characteristics of information flow and not the semantic properties. He is also not explicitly trying to use information theory to account for the phenomenal properties of consciousness—which would come later from Dretske (1995) and Tye (1995). Still Sayre does think his account will yield the "look" of a penny. He continues: "When...retinal information that has fed one's perception of a circular penny undergoes a major shift in grouping because of a major change in viewing angle, a new configuration will become stabilized on level C and the penny will assume a more ellipitical appearance" (Sayre, 1976, p. 151). His account would attempt to derive the phenomenal look of the penny (circular vs. elliptical) from the correlation through the informational cascade from penny to level B and on to C. The informational cascade would be analyzed in terms of "mutual information" available along the way in the informational channels of the visual system. Using informational concepts, Sayre develops a view that he calls "patterned visual response....a pattern of neuronal activity in the cortex is a set of events so ordered that the occurrence of a subset increases the probability that the remainder will occur in proportion to the size of the given subset. When a pattern of activity has been established in the cortex by repeated configurations on the retinal level, the occurrence of only part of this pattern during subsequent moments may stimulate the remainder into an active state" (Sayre, 1976, p. 152). Sayre sees the perceptual process as the development of these more or less stable patterned responses that the brain uses, along with stored information from memory, to guide an organism's behavior. He sees the whole process as constantly in flux. "Even when aware of an entirely common and unchanging perceptual object, a person's afferent information processing system is busily engaged in the endless activity of balancing stable neuronal structures against the diversity of information from his sensory receptors" (Sayre, 1976, p. 154). Sayre finds this account to be particularly good at explaining *constancy* phenomenon and the phenomenon of *filling-in*, since the patterned responses achieved can be maintained despite variation (or absences) of corresponding afferent input over time (Sayre, 1976, p. 157). A penny continues to look round when at an angle, occluded, or in one's blind spot, because the perceptual patterned response has been activated and is maintained, once acquired.

Is Sayre's account of perceptual consciousness an internalist or externalist view of the phenomenal content of perceptual consciousness? This sort of question arises in the recent works of Tye (1995) and Dretske (1995), but it is not clear that Sayre is addressing this issue in the same terms as recent discussion. It is clear that he accepts something of an Aristotelian account upon which there are structures of properties in external objects and those very patterns of structure can be conveyed into the perceiving mind in the form of information. He claims that "...it is possible to maintain with Aristotle that these objective structures in the domain of the mind's activity are identical with objective structures at the other end of the perceptual cascade. In a literal sense [made precise only in the context of information theory], structures present in the organism's cortex are identical with structures characterizing the object of perception" (Sayre, 1976, p. 155). This example from Aristotle suggests a very literal interpretation of *in-form-ation*. Sayre thinks that information and the concepts of mutual information in a noiseless channel supply the ingredients to make this clear. Somewhat surprisingly, Sayre does not employ the concept of a representation. He does not say that the mind represents the external structure of the perceived object or its properties. He seems to be saying, with Aristotle, that the mind acquires the form or structure of the external properties (perhaps setting up a sort of isomorphism of mind and object). At any rate, I would say that Sayre definitely falls on the externalist side of the qualia issue. That is, the phenomenal content of conscious perception derives from the external properties of objects (and indeed, seems to share something identical with them on his view).

When Sayre turns his attention to intention and purpose, he defines terms and then applies informational concepts. "A purpose is a neuronal configuration capable of shaping behavior in repeatable sequences, and capable on occasion of bringing that behavior under conscious control" (Sayre, 1976, p. 175). Sayre sees purposive behavior as terminating (when not frustrated) in a conscious patterned response of the perceived desired end-state. So his view of conscious perception contributes smoothly to his view of intentional or purposive action. Of course, he needs some account of how these purposes are acquired and of how they lead to the realization of the state of affairs that generates the patterned response at termination. Presumably, the latter bit is supplied by feedback control and the former by some process of concept formation (or something similar).[18]

When Sayre considers information and meaning, he makes the usual distinctions between "intension" and "extension" or reference, and identifies shared intension of a language community with something like his shared patterned response from perception. "The sense in which the same meaning may be present in the minds of different individuals is closely analogous to the sense in which the same informational structures may be present in the mind of the perceiver as are present in the object of veridical perception" (Sayre, 1976, p. 201). Sayre observes that "meanings are neuronal patterns removed from exclusive control by external objects, and brought under the control of verbal signals" (Sayre, 1976, p. 210). Sayre thinks this is important because "we have finally arrived at an understanding of language as a means of conveying information in the sense of meanings that can be communicated linguistically. The gap has been closed between 'information' of communication theory and 'information' in the sense of what is semantically meaningful" (Sayre, 1976, p. 201). Of course, he may have been overly optimistic, for without an explanation of false representational content it appears that his account is much closer to that of mutual informational content than to semantic content (though, Sayre would no doubt disagree). However, although he does discuss the important feature of meaning that it must become "dissociated" from the immediate environment, he does not offer an account of *falsity* (the real stumbling block for naturalized accounts of meaning).[19] Sayre actually has a view of concepts such that the conceptual content of a neuronal structure is tied to the perceptual state it would bring about if behavior it is driving were to be successful. So if one wants a banana, and it drives one's behavior successfully, one's behavior would culminate in the perception of a banana (Sayre, 1976, p. 217). One might call such a state a "virtual representation." Still one needs an account of what ties such a specific representation to its specific virtual content (prior to success and actual content), and Sayre does not supply this. I'm sure he would say it is via a perceptual process of conceptual abstraction, but the devil is in the details. What Sayre does say looks to be a holistic account of concepts (Sayre, 1976, p. 219), but his account just scratches the surface in this work.

Adams (1979) first applied an informational account to teleological functions. His basic view is that an item has a teleological function if it contributes to the goal of a goal-directed system via an information feedback chain of sustained causal dependency. Adams and Enc (1988) showed how the account blocks attributions of functions to events of accidental good fortune, and Enc and Adams (1992) displayed the advantages of this type of account over propensity accounts of functions. Since this account tied teleological functions to goal-directed systems, a cybernetic account of goal-directed systems followed (Adams, 1982). The account began along the lines of Rosenblueth et al. (1943), but departed from that account by adding the notion of a goal-state representation.[20] The representation determines the goal of a goal-directed system and then plays a role in goal-state comparison (minimizing the difference between the system's actual state and its goal-state). To some, this addition of a goal-state representation to the informational accounts of purposive systems was a much-needed improvement (Nissen, 1997, Chapter 2), for now systems could be directed toward states that may not now (or ever) exist. Behavior is goal-directed if it is the product of a goal-state representation, a process of negative feedback correction for error, and a causal dependency of the latter upon the former. Of course, to provide a naturalized account of purposive systems, one still needs to say how a goal-state representation arises out of purely natural causes. For this, the information-based accounts of representation and meaning provided the missing pieces. In addition, based upon this information-based account of goal-directed system, Adams extended his account to cover an information-based account of intentional action and solutions to problems of causal deviance, among others (Adams, 1986a, b, 1989, 1997; Adams and Mele, 1988, 1992).

4. The Jump from Information to Meaning

> *C* acquires its semantics, a genuine meaning, at the very moment when a component of its natural meaning (the fact that it indicates [or carries information about $_{\text{FA}}$] *F*) acquires an explanatory relevance. (Dretske, 1988, p. 84)

Dretske (1981, 1983, 1986, 1988) was perhaps the first philosopher to attempt to meet in a naturalized way Bar-Hillel's challenge of uniting the mathematical theory of information with a semantics.[21] It is not uncommon to think that information is a commodity generated by things with minds. Let's say that a naturalized account puts matters the other way around, viz. it says that minds are things that come into being by purely natural causal means of exploiting[22] the information in their environments. That is the approach taken by Dretske as he tried consciously to unite the cognitive sciences around the well-understood mathematical theory of communication.

Up to this time, nearly everyone realized that information and mathematical properties of informational amounts and their transmission were not the same thing as semantic content or meaning. There were accounts that attempted to bridge the gap, but perhaps Dretske's insight was to see clearly (more clearly than most) what needed to be done to accomplish this. In particular for something "F" to mean *F* (or that x is *F*), all other information carried in a signal (information about more than *F*) needed to be overshadowed, and, it must be possible for a signal to say that something is *F*, even when it is not. To accomplish this requires many subtasks along the way: tasks such as explaining how the distal rather than proximal informational source can be isolated or how one piece of information can be featured as the semantic content of a representation even though other pieces of information are carried simultaneously.

Information and truth: First, one has to shift the focus from average amounts of information carried by a communication channel and mathematical properties of such channels, to the informational value carried by a single signal about events at an informational source. This is sometimes called the "surprisal" value of a signal. Second, one has to identify the semantic content with the informational value of that single signal. Upon doing so, information becomes wedded to truth. "What information a signal carries is what it is capable of 'telling' us, telling us *truly*, about another state of affairs....information is that commodity capable of yielding knowledge, and what information a signal carries is what we can learn from it" (Dretske, 1981, p. 44). "Information is what is capable of yielding knowledge, and since knowledge requires truth, information requires it also" (Dretske, 1981, p. 45). Now here Dretske is talking about information that p where a signal that carries this information carries it unequivocally. This information will be carried when the probability that p, given the signal *s* is one (unity). Dretske is quick to point out that these probabilities are objective and mind independent such that "the amount of information contained in the signal depends, not on the conditional probabilities that *we can independently verify*, but on the conditional probabilities themselves" (Dretske, 1981, p. 56).

Information flows: after adding the "xerox principle" (Dretske, 1981, p. 57) requiring that if a signal A carries the information that B and B the information that C, A carries the information that C, Dretske maintains that "[w]hat communication theory (together...with the xerox principle) tells us is that for communication of *content*, for the transmission of a *message*, ...one needs *all* the information associated with that content" (Dretske, 1981, p. 60). That is, if a signal carries the information that *s* is *F*, it must carry "as much information about *s* as would be generated by *s*'s being *F*" (Dretske, 1981, p. 63).

The intentionality of information: information itself is an intentional commodity. Dretske makes this connection by showing that even amounts of information, if transmitted from one location to another, depend on conditional nomic dependencies of probabilities. Basically, the flow of information depends on law-like connections and *laws have a modal status* (and support counterfactuals about relations between properties). Hence, at least part of the answer to how "amounts" of information can be relevant to "contents" of signals is that even amounts of information are conveyed *by nomic dependencies between specific properties*. Putting these pieces together yields a definition of *informational content*: A signal *r* carries the information that *s* is *F* = The conditional probability of *s*'s being *F*, given r (and *k*), is 1 (but, given *k* alone, less than 1) (Dretske, 1981, p. 65).[23] Of course, these dependencies between properties are not (or seldom) one to one. So even if there are nomic conditional probabilistic dependencies — there is still work to be done getting from information to a signal's univocally meaning that *s is F*.

Nested information: the information that *t* is *G* is *nested* in *s*'s being *F* = *s*'s being F carries the information that *t* is *g* (Dretske, 1981, p. 71). This happens when the conditional probability that *t* is *G*, given *r* is one and the conditional probability that *s* is *F*, given *r*, is also one. Being *G* and *F* may be tied analytically or nomologically such that it is not possible for the one property to be instantiated without the other. Then the property dependencies are such that a signal that carries one piece of information will carry other pieces. Hence, given the nature of the world and nomic dependencies of properties, there is no such thing as *the* informational content of a signal or message. Any signal carrying the information that p necessarily carries the information that p or q, but no signal that means p necessarily means that p or q (the signal that means that p or q is "p or q"). Clearly there is such a thing as *the* semantic content of a signal or message. So the trick yet to be turned is to go from multiple pieces of information contained in a signal to a univocal semantic content. Turning this trick will simultaneously explain how a signal could be false, as it turns out. For in both cases there has to be a locking of unique and univocal semantic content to an informational signal or message (a symbol or signal "*s* is *F*" must lock to its meaning, to *s's being F*).[24]

In Dretske's (1981) account there were still three more important ingredients in the transition from information to meaning: primary representational dependency, digital coding, and the learning period (the last of which was abandoned in his subsequent account (Dretske, 1988)).[25] Let's start with primary representational dependency (a notion not far from what Fodor (1990) would later call "asymmetrical causal dependency"). When we see through a causal chain of events during visual perception, how is it that we see the distal cause of our perceptual experience, not the proximal cause (or any cause in between) when we do see the distal cause? And unless this can be done in perception, how could thoughts be about distal, not proximal objects, as well? Dretske's answer has two parts. First, if there are multiple pathways from the distal object to the percept, then when one perceives the external object, it is the distal object alone about which nondisjunctive information is carried. Second, if there are not multiple pathways from distal object to percept, then one may still perceive the external object because of the physical vehicles of the perceptual processes and mechanisms. Information about any more proximal causes may be at most nested in information carried about the more distal object. I hear the doorbell ringing, not my eardrum vibrating, even though information about the latter carries information about the former. Why then do I hear the latter? It is because the latter is represented in a *primary way* by the physical nature of my auditory system. "S gives *primary representation* to property B (relative to property g) = S's representation of something's being g depends on the informational relationship between B and g, but *not vice versa*" (Dretske, 1981, p. 160). (Those familiar with Fodor's new theory of meaning will see the similarity.) This basically says that my auditory percept is asymmetrically causally dependent upon information about the ringing bell, not the vibrating eardrum, though it carries both pieces of information. Primary representation is used to gets us representations of distal objects and properties. We can see them and think about distal objects without being blocked by causal (informational) intermediaries.

Still what about meaning? A thought about the ringing doorbell had better not *mean* something about my vibrating eardrum, regardless of whether it only carries that information in a non-primary way. And if semantic content has an informational origin, my thought would have both pieces of information in its origin and more. So how could my thought ever have univocal semantic content? How could my thought about the bell just mean *bell*, for example, since the ringing never carries information just about the bell? The answer again is nesting. Dretske calls the specific kind of coding that we need *digital* coding (Dretske, 1981, p. 185). A signal carries information in a *digital* way when all other information it may carry asymmetrically depends on the featured piece of information (again, familiar Fodorian themes may help to understand the basic idea).

We are still not to the level of meaning, even if a structure carries a unique piece of information in a "digital way." For *information cannot be false,* but I can falsely think or say that the doorbell is ringing when it is not. So the final step from information to meaning is to kick away the ladder — to explain how meaning, while depending on informational origins, can abandon them when something acquires a full semantic content (i.e., meaning). For this, Dretske (1981, p. 193) appeals to a *learning period L*.[26] "Suppose...that during *L* the system develops a way of digitalizing the information that something is *F*: a certain type of internal state evolves which is selectively sensitive to the information that *s* is *F*....Once this structure is developed, it acquires a life of its own, so to speak, and is capable of conferring on its subsequent tokens...its semantic content (the content it acquired during *L*) *whether or not those subsequent tokens actually have this as their informational content*." In short, the structure type acquires its meaning from the sort of information that led to its development as a cognitive structure....What this means, of course, is that subsequent tokens of this structure can *mean* that *s* is *F*, can have this propositional content, despite the fact that they fail to carry this information, despite the fact that the *s* (which triggers their occurrence) is not *F*" (Dretske, 1981, p. 193). "We have...meaning without truth" (Dretske, 1981, p. 195).

Dretske (1988) later abandoned the notion of a learning period and replaced it with an explanatory role account. His view now is that a cognitive structure "F" means *F* (that something is F), when "F"'s indicate *F*s (carry information about *F*s) and their so doing explains relevant bodily movements *because* they indicate *F*s.[27]

Fodor (1987, 1990, 1994, 1998) like Dretske saw what was required to make the jump form information to meaning. Indeed, Fodor (1990) was quick to point out the failings in Dretske's (1981) first attempt to naturalize meaning, as is well known.[28] Fodor's own *asymmetrical causal dependency* theory of meaning has many similarities to (and yet many differences from) Dretske's account. It is similar in that Dretske's notions of *primary representation* and *nesting* are types of asymmetrical dependencies. Its chief difference is that no part of Fodor's account depends upon the notion of information or indication or any concept directly derived from the origins of the mathematical theory of information.

On Fodor's view, a cognitive structure "X" means X when, it's a law that 'Xs cause "X"s' and when, for anything Y that causes "X"s, Ys would not cause "X"s, but for the fact that Xs do. A false tokening of an "X" is produced by a Y or a Z, say, because "X"s are dedicated (locked) to Xs by the asymmetrical dependency of laws. So a tokening of an "X" by a Y constitutes a misrepresentation of a Y as an X. Now occasionally Fodor refers to his theory as a "pure" informational theory of semantic content (Fodor, 1990, 1994), but as near as I can tell, Fodor's theory is a theory of laws not of information. In some environments, there may be information relations between A-events and B-events, but not be a law of nature that ranges over A-events and B-events. For instance, if I am perfectly reliable and honest, if I say "B" then B, but there is hardly a law of nature that if Adams says "B" then B. Under these conditions I am a reliable conduit of information — the information that B. So while there is certainly an historical connection between the informational turn and Fodor's theory of meaning, I will not recount the details of Fodor's theory of meaning here, for I do not think it is an information-based account, despite Fodor's affinity for calling it one. In any case, I have said more than enough about Fodor's theory of meaning elsewhere (Adams, 2003b).

Perry et al. (1983, 1985, 1990, 1991): In a series of papers and books (Barwise[29] and Perry, 1983, 1985; Perry and Israel, 1990, 1991), Perry and co-workers developed an account of information that was intended to apply primarily to situation types (states of affairs). The account bears some similarities to that of Dretske (1981), yet there are differences (for example, theirs does not obviously employ the mathematical theory of information). The account of Perry and co-workers bears some similarity to that of Carnap and Bar-Hillel[30] (1964), though there are differences here as well (Perry and co-workers do not intend their account to flow from or support an inductive theory of probability, as was true of Carnap and Bar-Hillel).

Anyone familiar with Dretske's account of Shannon-Waver information first (and Perry et al., second) will be taken aback by reading Perry et al. The terminology is completely different. One needs a translation manual. So I shall supply the beginnings of that here, but I shall not try to give an exhaustive key to moving back and forth between the two accounts.[31] Central notions in the account of Perry and co-workers are *constraint* and *involvement*. Constraints in the Perry account play the role of *channel conditions* in the Dretske account. Dretske (1981, p. 115) defines channel conditions in this way: "The channel of communication = that set of existing conditions (on which the signal depends) that either (1) generate no (relevant) information, or (2) generate only redundant information (from the point of view of the receiver)." These are stable background conditions, whose stability, allow something (say, a voltmeter) to carry information about something else (voltage of a battery). "If initial checks determine that the spring [in the voltmeter] has not broken, worked loose, or lost its elasticity, then the integrity of the spring qualifies as a channel condition for (at least) hours, probably days, and perhaps event months... (Dretske, 1981, p. 118). If channel conditions are locally stable, the voltmeter reading carries information about the battery ("12" carries information that the battery is a fully charged 12-V battery, for the conditional probability of he latter, given the former is one).

Perry et al. would say, instead, that there are "constraints" in place such that the voltmeter reading *involves* the 12-V drop across the battery leads, but this comes to much the same thing as Dretske's channel conditions (as nearly as I can tell). All of the examples of Perry et al. are of complete, unequivocal information exchange (Dretske's conditional probabilities of 1). For that reason, Perry et al. employ no formulae for calculating probabilities and amounts of information less than full information contained in one event about another event (hence, one finds no calculations of bits and bytes).

Here is a typical passage putting all of these notions together: "Suppose we have a constraint that a thing of type T is also of type T'. Relative to that constraint, the fact that x is T indicates that x is T'. I'll call these the indicating fact and the indicated proposition. A given tree has one hundred rings. That indicates that it is at least one hundred years old" (Perry, 1990, p. 177).

Importantly, Perry develops his account of information to help him solve familiar puzzles of intentionality (Perry, 1990). He and his co-authors explicitly develop informational notions designed to apply to puzzles about belief and belief ascription. For example, they even adopt the relevant terminology of featuring informational "modes of presentation" (Perry and Isreal, 1991). Since informational relations largely involve nomic relations of dependency between *types* of events, Perry realizes that it will be a struggle to adapt informational concepts to apply to *individuals*. I think that, if there is a single most distinctive feature of Perry's account (in contrast to Dretske's), it is Perry's application to belief contexts about individuals. He tries to use informational concepts to explain how one can have the belief that *a* is *F* (not just explaining how something could mean *F*) and to relate information to *modes of presentation* (in part to explain in informational terms how S could believe that *a* is *F*, but not that *b* is *F*, when *a* = *b*).

Let me follow another of the key examples used to make this transition (Perry, 1990; Perry and Israel, 1990). Suppose the veterinarian takes an X-ray of Jackie (the Barwises' dog). Jackie has a broken leg. Patterns on the exposed X-ray carry information about the broken leg, but how do we go from the fact that they carry information that *an animal has a broken leg* ("pure information") to the information that it is *Jackie* who has the broken leg ("incremental information")? That is the relevant question to which Perry wants to apply informational concepts.

Perry and co-workers argue that there are "three interconnected types:"
T: x exhibits such and such a pattern
T': x was an X-ray of y
T": y has a broken leg (Perry, 1990, p. 178)

Constraint 1: "T involves T' & T". If an X-ray looks like this [demonstrative indicating the pattern on the X-ray], there was an animal it was taken of, who had a broken leg." Now add "Given that the X-ray was taken of Jackie, it seems that the X-ray's exhibiting such and such pattern indicates that Jackie has a broken leg" (Perry, 1990, p. 179).

As Perry says, the vet has a mode of presentation "being the dog X-rayed" of Jackie. This mode is relevant in connecting the X-ray with Jackie. When the vet says that the X-ray shows that Jackie has a broken leg, this is only true, given that it was an X-ray *of Jackie* (and that it is a recent X-ray of Jackie). Of course, the X-ray *says* none this as a matter of its "pure" information, but only as a matter of its "incremental" information. When that is the case, the mode of presentation "becomes essentially irrelevant," and the vet essentially sees the X-ray as containing information *about Jackie now*.

Interestingly, Perry goes on to use these informational concepts to reconcile the supposed differences between the Russellians and Fregeans on matters of direct reference, modes of presentation, and the occurrence of individuals in propositions. This is a use to which informational concepts had not previously been put.

Perry and Isreal (1991) also are interested in showing how information can be relevant and used in a system. They are interested in showing how syntactic engines respect semantics, to use a familiar phrase. For this they introduce what they call *architectural constraints*. Again it involves coordination between modes of presentation and the information contained in those modes (as vehicles of content). Here is an example. Consider the relationship between a scale that weighs and measures the height of the same person. "...if a weight bar and height bar are connected that way [such that height and weight bars measure a single person at a specific geographic location and relation to the apparatus] the person whose head contacts the height bar is the person who is affecting the weight bar. We call the sort of constraint involved...an *architectural constraint* and the relation between subject matters (in our case, identity), the *architecturally grounded relation*. Information relative to architectural conditions and constraints, we call *architectural*" (Perry and Isreal, 1991, p. 150). In cases of human apparatus, we arrange the architectural constraints so that they give us both pure information and incremental information about the individual currently standing on the scale. In computers, we build them so that due to architectural constraints of the computer engineers, the syntax (and what it causes) respects the semantics. Of course, in the mind on the hoof as it were, nature has to take the role of the computer engineer. As Dennett (1976) long has pointed out from his "intentional stance," natural selection provides the background for the architectural constraints (if only Dennett were a realist about content). At any rate, not only are Perry and co. interested in saying how information can be about individuals, but also how information can be put to explanatory work.[32]

5. Conclusion

> In the beginning there was information. The word came later. The transition was achieved by the development of organisms with the capacity for selectively exploiting this information in order to survive and perpetuate their kind.
>
> (Dretske, 1981, p. vii)

From its very inception as a theoretical entity, *information* has been seen as a key ingredient in the making of a mind. Minds are usually associated with goal-directed or purposive behavior. So it is no surprise that immediately, Wiener (1943, 1948) developed his new *cybernetics* as a science of information guiding purposive behavior. And as we have seen many of the philosophers in the middle portions of the second half of the 20th Century followed this lead. As we have also seen, to fully implement the application of information to goal-directed (purposive) behavior, there needed to be an account of representation that could sustain the naturalization of teleology. A system or organism had to be able to acquire a mind (or at least acquire cognitive states that were capable of determining goal-states that may or may not exist at the time a system's behavior begins or that could be falsely tokened). For this to be possible, there had to be an account of cognitive representation capable of supporting goal-directed and even high level intentional behavior and semantic content.

Naturalizing the mind would require that purely natural (physical) causes be capable of being ingredients in the production of a mind (and mental representation). The main ingredient chosen for developing a theory of mind and representation, in the partial history I just told, has been *information*. Those who take the informational turn see information as the basic ingredient in building a mind. Information has to contribute to the origin of the mental. From informational beginnings, minds have to be able to represent (semantically mean) types of states of affairs. At a minimum, a goal-directed system has to have representations of the types of states of affairs it seeks to bring into existence (its goal-states). Thus, something in the system has to semantically represent those types of states. The system must also be able to represent its actual states (and the state of its environment), and then compare and minimize the differences between goal-state and actual state.

Collectively the philosophers above saw what needed to be done and saw that information was a key ingredient in understanding how purposive systems work. Different philosophers saw different pieces of the overall picture, and contributed and moved the project forward in different ways and to different degrees. There is still much more to be done[33], but there is no turning back. Like the Hotel California, once you take the informational turn — you can check in (you even can check out), but you can never leave.

Notes

[1]Someone someday should write a more inclusive treatment of the informational turn, looking at all of philosophy and much of computation theory and artificial intelligence research. I am sure that I am not the one to do this, and I know that this is not the time, but it would be a welcome addition to the literature. Sadly, my treatment here will be quite narrowly confined. I would like to thank Dan Dennett, Fred Dretske, Kenneth Sayre, and especially Luciano Floridi for helpful correspondence in preparation of this paper. Robin Andreasen and Paul Skokowski gave helpful comments on earlier drafts. I would also like to thank Tim Jones for helping me find the sources I needed. Finally, I would like to dedicate this paper to the memory of Berent Enc.

[2]Bar-Hillel puts the start at between 1928 and 1948, referring to the early work of Hartley (1928). He properly credits Hartley with christening the formula used to measure the amount of information generated by signal *s,* viz. log to the base 2 of 1/probability of the event *s* (among equally probable events). MacKay (1969) puts the start of the information age at 1950 too and also provides a useful history, describing the influence of Denis Gabor, and R.A. Fisher on the foundations of information theory.

[3]It is interesting to note that both Wiener and Turing were students of Bertrand Russell. Indeed, in the late 1940s and early 1950s it was a small world indeed, as can be seen in the list of credits given by these early information-theoretical pioneers. Wiener notes his indebtedness to Walter Cannon, Arturo Rosenblueth, and Julien Bigelow and John von Neumann. Wiener also notes that his efforts to understand the nature of controlled or purposive behavior required a method of calculating the amount of information generated by an event or signal. R. A. Fisher was working on the problem from the perspective of classical statistical theory, Shannon from the problem of coding information, and Wiener himself from the perspective of problems of noise and electrical filters. Wiener also notes that Kolmogoroff in Russia had done work similar to that of Wiener, simultaneously. Wiener also notes the influence upon the field of his colleagues Warren McCulloch and Walter Pitts, and Rashevsky, who were interested, as was Turing, in a logical calculus that could be implemented in a nervous system.

[4]One can begin to get a glimpse of the tight circle of folks all working on similar themes across different domains of study at the midpoint of the 20th Century by reading the introduction to Wiener's *Cybernetics*. The beginnings of the information age were clearly exciting times and formed the foundations of the current cognitive sciences.

[5]In this paper I must refuse to discuss a large portion of the literature. One way (not entirely happy) to narrow that literature is to not discuss the many papers that are critical discussions of some of the main attempts to apply information theory to philosophy. So, for example, sadly I will not discuss the many objections to Dennett, Dretske, Sayre, and others (including my own work). This strategy does have its advantages and disadvantages, but I see no other way to proceed in the space allowed.

[6]I discuss information theory as applied to knowledge in more depth elsewhere (Adams, 2003a).

[7]For a survey of the types of problems Dretske thought he could avoid by using information theory, see Adams and Clarke (manuscript) where we chronicle them in great detail.

[8]Of course, this was not the case in the non-philosophy departments in the cognitive sciences. Information processing models in psychology were going strong by the 1970s, for instance. Some areas such as the study of signal detection were very quick to adopt informational concepts.

[9]As everyone knows, it is no accident that the informational age and the wealth of activity surrounding the development of mathematical theories of information and computation came during a war. Intellectual interest often follows the flow of funds, and the funds definitely flowed into projects surrounding computing and cybernetic systems for purpose of war (sadly). Wiener's early work had application to targets, missiles, and the control systems leading up to the so-called "smart bombs" of today.

[10]While this is true of what Hartley was trying to capture in his analysis of information, Bar-Hillel cannot resist pointing out that apparently even Hartley was sometimes taken in by the ambiguity of the quantitative versus the semantic use of "information" (Bar-Hillel, 1955, p. 95).

[11]I was fortunate to be at Wisconsin during the rise of what Fodor (1990) dubbed "Wisconsin Semantics." While Dretske (1981) was working on adapting informational concepts to knowledge and meaning, Enc (1979) and I (1979) were struggling to adapt such concepts to problems of goal-directed systems and teleological functions. Of course, Stampe (1975) had influenced the group with his causal theory of meaning, and Sober (1985) was working on applications to problems in the philosophy of biology. These were exciting times, indeed.

[12]For more on my own reconstruction of attempts to naturalize meaning along these lines see (Adams, 2003b).

[13]In fairness to MacKay, in a 1969 "postscript" to his orginal text of the 1950s he added comments about his "operational" notion of information. "It should perhaps be emphasised that the meaning of a message is not here defined as the effect it produces in the recipient. The meaning it has for a given recipient is a logical property, which it possesses whether or not it has a chance to exercise its function on the recipient. The definition of meaning in the original paper as the 'selective operation which the signal performs on the set of possible states of readiness' was open to misinterpretation on this score, since it was easy to confuse the logical operation on the set with the physical operation on the organism" (MacKay, 1969, pp. 77–78).

[14]Wooldridge knew of the early work on the neural circuits (Lettvin et al., 1959) of the frog and was plugged into the tight circle of people interested in cybernetics mentioned by Wiener (1948) in his introduction.

[15]When first I asked Dennett how he came to look to information theory for help, he replied "that's a very good question." He did not immediately mention the work of MacKay, but MacKay gets six index references in Dennett's (1969) book, and Shannon and Wiener get none. I notice that Dennett (1994) also mentions the influence of MacKay in his "Self-Portrait." Dennett and MacKay did in fact meet at least twice later in life, but long after 1969.

[16]I first met Dennett in 1978 at Andrew Woodfield's year-long Philosophy of Mind Workshop at the University of Bristol. I had come from Wisconsin and Dennett was very interested to hear about the prospects of using informational concepts to understand semantic content. Dennett was a friend of Stampe's from graduate school at Oxford and knew that Stampe (1975) had been working on a causal theory of semantic content. Though Dennett was acquainted with the mathematical theory, given his own view of that theory, he must have been suspicious of the outcome from the start. Nonetheless, he was interested in learning details of how one would go about attempting to ground cognitive content in information.

[17]Sayre's 1986 BBS target article (and peer commentary) updates his views, but I still think his 1976 book best illustrates his contribution to the informational turn. The 1986 BBS article still lacks an informational account of semantic content that can be falsely tokened, and many peer commentators take him to task.

[18]Sayre is surprisingly silent on the matters of goal-state formation and feedback control in the section of the book where he covers purpose and intention (Sayre, 1976, pp. 172–177). He says that to be motivated by the purpose of finding one's pencil is to be guided by neuronal activities that terminate in the perceiving of the pencil. But that doesn't explain what makes the perception the goal. Death terminates many activities too, but is not necessarily one's purpose or goal. So there is a gap that needs filling to explain how finding the pencil becomes the purpose and what determines or fixes that as the content of the goal. This Sayre does not attempt to supply.

[19]Note Sayre's appeal to the analogy with "veridical" perception.

[20]Though he is not a philosopher, I found the book by Powers (1973) to be particularly influential in my own thinking.

[21]Bar-Hillel's (1964) own semantic account of information was far from a naturalized account.

[22]Exploiting in a non-intentional sense.

[23]While Dennett (1969, 1987) correctly pointed out that the information one receives depends upon what one already knows, Dennett went on to conclude from this that information is thereby "not amenable to precise quantification." Dretske (1981, p. 79) maintains to the contrary "unlike Dennett, I do not believe this means that we cannot precisely quantify the amount of information contained in a signal." That something is a relative quantity does not mean it is not precisely measurable from a frame of reference (consider space-time position or acceleration...relative, yet quantifiable) (Adams, 2002). And, as Dretske goes on to point out, very often the relevant background knowledge of various agents is equal and any differences drop out. In the limit, no background knowledge ($k = 0$), the amount of information generated by an event reverts to an absolute quantity.

[24]However, perhaps somewhat surprisingly, the possibility of falsity brings with it potential loss of information. Signals that carry the information that *s* is *F* cannot be false. Signals that mean that *s* is *F* can be false. Hence, meaning is not information — though it may have informational origins, and will have, if this account is correct.

[25]Dretske replaces the "learning period" of his 1981 account with an "indicator-causal" explanatory account in 1988, Chapter 4. He flirted with a different account of meaning (and misrepresentation in his 1986), but abandoned that for the 1988 account.

[26]This move was later successfully criticized by Fodor (1990) and replaced by Dretske's new account of the last step (Dretske, 1988, Chapter 4). So it took Dretske three tries to get a stable account of meaning (Dretske, 1981, 1986, 1988). Many people mistakenly think his second attempt (Dretske, 1986) was his final view.

[27]See Adams (2003b) for more about Dretske's current view and how he got there.

[28]See Fodor's (1990) "Semantics, Wisconsin Style," first published in 1984.

[29]Barwise was in the math department at Wisconsin while I was there. Later he moved to Stanford, teaming up with Jon Perry. Barwise attended several philosophy department lectures by Dretske about information and cognition in the late 1970s. I didn't meet Perry and Isreal until I was at Stanford's C.S.L.I. in 1992.

[30]See Carnap and Bar-Hillel's "An Outline of a Theory of Semantic Information" reprinted from 1952, in Bar-Hillel's 1964 book.

[31]There is even some questions by Dretske about whether that can be done fully (Dretske, 1985).

[32]See Fodor (1986) for dissent on the Barwise/Perry account and Dretske (1990) for an account of putting information to work that requires learning.

[33]We still need an account of how mental states cause behavior in virtue of their representational content, and how there can be psychological laws that quantify over mental states in virtue of their content. This is not to say that progress in this regard has not been made. It has. We took a wrong turn in the battle between narrow vs. broad content, but now seem to be back on track. Still there is much left to be done, having taken the informational turn heading toward naturalizing the mind.

References

Adams, F. (1979), ‘A Goal-State Theory of Function Attributions’, *Canadian Journal of Philosophy* 46, pp. 498–518.

Adams, F. (1982), *Goal-Directed Systems*, Ann Arbor: University Microfilms International.

Adams, F. (1986a), ‘Intention and Intentional Action: The Simple View’, *Mind & Language* 1, pp. 281–301.

Adams, F. (1986b), ‘Feedback About Feedback: Reply to Ehring’, *Southern Journal of Philosophy* 24, pp. 123–131.

Adams, F. (1989), ‘Tertiary Waywardness Tamed’, *Critica* 21, pp. 117–125.

Adams, F. (1997), ‘Cognitive Trying’, in G. Holmstrom-Hintikka & R. Tuomela, eds., *Contemporary Action Theory*, Vol. I, Dordrecht: Kluwer, pp. 287–314.

Adams, F., manuscript, Informational Indeterminacy, Invited Commentary on “A Deeper Problem for Dretske’s Theory of Information Content” (by Andrea Scarantino, University of Pittsburgh), Society for Philosophy & Psychology, Edmonton, Alberta, June 2002

Adams, F. (2003a), ‘Knowledge’, in L. Foridi, ed., *The Blackwell Guide to the Philosophy of Information and Computing*, Oxford: Basil Blackwell, Chapter 7.

Adams, F. (2003b), ‘Thoughts and their Contents: Naturalized Semantics’, in Warfield & Stich, eds., *The Blackwell Guide to Philosophy of Mind*, Oxford: Basil Blackwell, pp. 143–171.

Adams, F. and Clarke, M., manuscript, Resurrecting the Tracking Theories.

Adams, F. and Enc, B. (1988), ‘Not Quite By Accident’, *Dialogue*, 27, pp. 287–297.

Adams, F. and Mele, A. (1988), ‘The Role of Intention in Intentional Action’, *Canadian Journal of Philosophy* 19, pp. 511–532.

Adams, F. and Mele, A. (1992), ‘The Intention/Volition Debate’, *Canadian Journal of Philosophy* 22, pp. 323–338.

Armstrong, D. (1968), *A Materialists Theory of the Mind*, London: Routledge, & Kegan Paul.

Bar-Hillel, Y. (1955), ‘An Examination of Information Theory’, *Philosophy of Science* 22, pp. 86–105.

Bar-Hillel, Y. (1964), *Language and Information*, Reading, MA: Addison-Wesley.

Barwise, J. (1986), ‘Information and Circumstance’, *Notre Dame Journal of Formal Logic*, 27.

Barwise, J. and Perry, J. (1983), *Situations and Attitudes*, Cambridge, MA: MIT/Bradford.

Barwise, J. and Perry, J. (1985), ‘Shifting Situations and Shaken Attitudes: An Interview with Barwise and Perry’, *Linguistics and Philosophy* 8, pp. 105–161.

Dennett, D. (1969), *Content and Consciousness*, London: Routledge & Kegan Paul.

Dennett, D. (1981), *Brainstorms*, Cambridge, MA: MIT/Bradford.

Dennett, D. (1986), ‘Engineering’s Baby’, *Behavioral and Brain Sciences* 9, pp. 141–142.

Dennett, D. (1987), *The Intentional Stance*, MA: MIT/Bradford.

Dennett, D. (1994), ‘Self-Portrait’, in Guttenplan, S., ed., *A Companion to the Philosophy of Mind*, Oxford: Blackwell, pp. 236–244.

Dretske, F. (1981), *Knowledge and the Flow of Information*, Cambridge, MA: MIT/Bradford.

Dretske, F. (1983), ‘Precis of Knowledge and the Flow of Information, *Behavioral and Brain Sciences* 6, pp. 53–56.

Dretske, F. (1985), ‘Constraints and Meaning’, *Linguistics and Philosophy* 8, pp. 9–12.

Dretske, F. (1986), ‘Misrepresentation’, in R. Bogdan, ed., *Belief*, Oxford: Oxford University Press.

Dretske, F. (1988), *Explaining Behavior*, Cambridge, MA: MIT/Bradford.

Dretske, F. (1990), ‘Putting Information to Work’, in P. Hanson, ed., *Information, Language, and Cognition*, Vancouver: University of British Columbia Press, pp. 112–124.

Dretske, F. (1991), ‘Replies’, in B. McLaughlin, ed., *Dretske and His Critics*, Oxford: Basil Blackwell.

Dretske, F. (1995), *Naturalizing the Mind*, Cambridge, MA: MIT/Bradford.

Enc, B. (1979), 'Function Attributions and Functional Explanations', *Philosophy of Science* 46, pp. 343–365.
Enc, B. and Adams, F. (1992), 'Functions and Goal-Directedness', *Philosophy of Science* 59, pp. 635–654 (reprinted in Allen, C., Bekoff, M. and Lauder, G., eds., *Nature's Purposes*, Cambridge, MA: MIT/Bradford).
Feigenbaum, E. and Feldman, J. 1961: *Computers and Thought*, New York: McGraw-Hill.
Fetzer, J. Forthcoming, 'Information, Misinformation, and Disinformation'.
Floridi, L. Forthcoming a, 'Outline of a Theory of Strongly Semantic Information', *Minds and Machines*.
Floridi, L. Forthcoming b, 'Is Semantic Information Meaningful Data?' *Philosophy and Phenomenological Research*.
Fodor, J. (1986), 'Information and Association', *Notre Dame Journal of Formal Logic*, 27.
Fodor, J. (1987), *Psychosemantics*, Cambridge, MA: MIT/Bradford.
Fodor, J. (1987), 'A Situated Grandmother? Some Remarks on Proposals by Barwise and Perry', *Mind and Language* 2, pp. 64–81.
Fodor, J. (1990), *A Theory of Content and Other Essays*, Cambridge, MA: MIT/Bradford.
Fodor, J. (1994), *The Elm and the Expert*, Cambridge, MA: MIT/Bradford.
Fodor, J. (1998), *Concepts*, Cambridge, MA: MIT/Bradford.
Gibson, J. (1979), *The Ecological Approach to Visual Perception*, Boston: Houghton Mifflin.
Grice, P. (1957), 'Meaning', *Philosophical Review* 66, pp. 377–388.
Hartley, R. (1928), 'Transmission of Information', *Bell Stystem Technical Journal* 7, pp. 535–563.
Hanson, P. (1990), *Information, Language, and Cognition*, Vancouver: University of British Columbia Press.
Lettvin, J., Maturana, H., McCulloch, W. and Pitts, W. (1959), 'What the Frog's Eye Tells the Frog's Brain', *Proceedings of the Institute of Radio Engineers* 47, pp. 1940–1951.
MacKay, D. (1951), 'Mindlike Behavior in Artefacts', *The British Journal for the Philosophy of Science* (reprinted in Sayre, K. and Crosson, F., eds., *The Modeling of Mind*, New York: Simon & Schuster, 1963).
MacKay, D. (1956), 'Towards an Information-Flow Model of Human Behaviour', *British Journal of Psychology* 43, pp. 30–43.
MacKay, D.M. (1969), *Information, Mechanism and Meaning*, Cambridge, MA: MIT Press.
Marr, D. (1982), *Vision*, New York: W. H. Freeman & Co.
McLaughlin, B. (1990), *Dretske and his Critics*, Oxford: Basil Blackwell.
Nissen, L. (1997), *Teleological Language in the Life Sciences*, Lanham, Md: Rowman & Littlefield.
Perry, J. (1990), 'Individuals in Informational and Intentional Content', in E. Villenueva, ed., *Information, Semantics and Epistemology*, Oxford: Basil Blackwell, pp. 172–189.
Perry, J. and Isreal, D. (1990), 'What is Information?' in P. Hanson, ed., *Information, Language, and Cognition*. (Volume 1, Vancouver Studies in Cognitive Science), Vancouver: University of British Columbia Press.
Perry, J. and Isreal, D. (1991), 'Information and Architecture', in J. Barwise, M. Gawron, G. Pltokin and S. Tutiya, eds., *Situation Theory and Its Architecture*, vol 2, Stanford: Stanford University.
Powers, W. (1973), *Behavior: The Control of Perception*, London: Wildwood House.
Rosenblueth, A, Wiener, N. and Bigelow, J. (1943), 'Behavior, Purpose & Teleology', *Philosophy of Science* 10, pp. 18–24.
Russell, B. (1921), *The Analysis of Mind*, London: Allen & Unwin, Ltd.
Sayre, K. and Crosson, F. (1963), *The Modeling of Mind: Computers and Intelligence*, New York: Simon & Schuster.
Sayre, K. (1965), *Recognition*, Notre Dame: University of Notre Dame Press.
Sayre, K. (1976), *Cybernetics and the Philosophy of Mind*, London: Routledge & Kegan Paul.
Sayre, K. (1969), *Consciousness*, New York: Random House.

Sayre, K. (1986), 'Intentionality and Information Processing: An Alternative Model for Cognitive Science', *Behavioral and Brain Sciences* 9, pp. 121–160.

Sayre, K. (1987), 'Cognitive Science and the Problem of Semantic Content', *Synthese* 70, pp. 247–269.

Shannon, C. and Weaver, W. (1949), *The Mathematical Theory of Communication*, Champaign: University of Illinois Press. (Reprint with a new introduction by Weaver, of Shannon's work by the same name in 1948 in the *Bell Systems Technical Journal*.)

Sober, E. (1985), *The Nature of Selection*, Cambridge, MA: MIT/Bradford.

Sommerhoff, G. (1974), *Logic of the Living Brain*, New York: John Wiley & Sons.

Stalnaker, R. (1984), *Inquiry*, Cambridge, MA: MIT/Bradford.

Stampe, D. (1975), 'Toward a Causal Theory of Linguistic Representation', in P. French et al., eds., *Midwest Studies in Philosophy*, 2, Minneapolis: University of Minnesota Press.

Turing, A. (1950), 'Computing Machinery and Intelligence', *Mind* 59, pp. 433–460.

Tye, M. (1995), *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*, Cambridge, MA: MIT/Bradford.

Weiner, N. (1948), *Cybernetics*, New York: John Wiley & Sons.

Wooldbridge, D. (1963), *The Machinery of the Brain*, New York: McGraw-Hill.

Young, P. (1987), *The Nature of Information*, New York: Praeger Publishers.