THE BEHAVIORAL AND BRAIN SCIENCES (1983) 6, 55–90
Printed in the United States of America

Précis of *Knowledge and the Flow of Information*

Fred I. Dretske
Department of Philosophy, University of Wisconsin, Madison, Wisc. 53706

Abstract: A theory of information is developed in which the informational content of a signal (structure, event) can be specified. This content is expressed by a sentence describing the condition at a source on which the properties of a signal depend in some lawful way. Information, as so defined, though perfectly objective, has the kind of semantic property (intentionality) that seems to be needed for an analysis of cognition. Perceptual knowledge is an information-dependent internal state with a content corresponding to the information producing it. This picture of knowledge captures most of what makes knowledge an important epistemological notion. It also avoids many of the problems infecting traditional justificational accounts of knowledge (knowledge as "justified, true belief"). Our information pickup systems are characterized in terms of the way they encode incoming information (perception) for further cognitive processing. Our perceptual experience is distinguished from our perceptual beliefs by the different way sensory information is encoded in these internal structures. Our propositional attitudes – those (unlike knowledge) having a content that can be either true or false (e.g., belief) – are described in terms of the way internal (presumably neural) structures acquire during learning a certain information-carrying role. The content of these structures (whether true or false) is identified with the kind of information they were developed to carry.

Keywords: belief; cognition; concept; information; intentionality; knowledge; meaning; perception; representation; semantics

Knowledge and the Flow of Information (Dretske 1981; henceforth *Knowledge*) is an attempt to develop a philosophically useful theory of information. To be philosophically useful the theory should: (1) preserve enough of our common understanding of information to justify calling it a theory *of* information; (2) make sense of (or explain its failure to make sense of) the theoretically central role information plays in the descriptive and explanatory efforts of cognitive scientists; and (3) deepen our understanding of the baffling place of mind, the chief consumer of information, in the natural order of things.

A secondary motive in writing this book, and in organizing its approach to philosophical problems around the notion of information, was to build a bridge, if only a terminological one, to cognitive science. Even if we don't have the same problems (psychologists are no more interested in Descartes's Demon than philosophers are in Purkinje's twilight shift), we have the same subject, and both sides could profit from improved communication.

In pursuit of these ends, it was found necessary to think of information as an *objective* commodity, as something whose existence (as information) is (largely) independent of the interpretative activities of conscious agents. It is common among cognitive scientists to regard information as a creation of the mind, as something we conscious agents assign to, or impose on, otherwise meaningless events. Information, like beauty, is in the eye of the beholder. For philosophical purposes though, this puts things exactly backward. It assumes what is to be explained. For we want to know what this interpretative ability amounts to, why some physical systems (typically, those with brains) have this capacity and others do not. What makes *some* processors of information (persons, but not television sets) sources of meaning? If we *begin* our study by populating the world with fully developed cognitive systems, systems that can transform "meaningless" stimuli into thoughts, beliefs, and knowledge (or whatever is involved in interpretation), we make the analysis of information more tractable, perhaps, but only by abandoning it as a tool in our quest to understand the nature of cognitive phenomena. We merely postpone the philosophical questions.

Part I of *Knowledge* develops a semantic theory of information, a theory of the propositional *content* of a signal (event, structure, or state of affairs). It begins by rehearsing some of the elementary ideas of the mathematical theory of communication (Shannon & Weaver 1949). This theory, though developed for quite different purposes, and though having (as a result) only the remotest connection (some would say *none*) with the kinds of cognitive issues of concern to this study, does, nonetheless, provide a key that can be used to articulate a semantical theory of information. Chapters 2 and 3 are devoted to *adapting* and *extending* this theory's account of an information source and channel into an account of how much information a *particular* signal carries about a source and what (if any) information this is.

Part II applies this theory of information to some traditional problems in epistemology: knowledge, skepticism, and perception. Knowledge is characterized as information-produced belief. Perception is a process in which incoming information is coded in analog form in preparation for further selective processing by cognitive (conceptual) centers. The difference between seeing a

 0140-525X/83/010055-36/$06.00

duck and recognizing it *as* a duck (seeing *that* it is a duck) is to be found in the different way information about the duck is coded (analog vs. digital).

Part III is devoted to an information–theoretic analysis of what has come to be called our propositional attitudes – in particular, the belief that something is so. Belief, the *thinking* that something is so, is characterized in terms of the instantiation of structures (presumably neural) that have, through learning, acquired a certain information-carrying role. Instances of these structures (the types of which are identified as concepts) sometimes fail to perform satisfactorily. This is false belief.

Information

The mathematical theory of communication (Cherry 1951; Shannon & Weaver 1949) is concerned with certain statistical quantities associated with "sources" and "channels." When a certain condition is realized at a source, and there are other possible conditions that might have been realized (each with its associated probability of occurring), the source can be thought of as a generator of information. The ensemble of possibilities has been reduced to a single reality, and the amount of information generated is a function of these possibilities and their associated probabilities. The die is cast. Any one of six faces might appear uppermost. A "3" appears. Six possibilities, all (let us say) equally likely, have been reduced to one. The source, in this case the throw of the die, generates 2.6 bits of information ($\log_2 6 = 2.6$).

But more important (for my purposes and for the purpose of understanding *communication*) is the measure of how much information is transmitted from one point to another, how much information there is at point r (receiver) about what is transpiring at s (source). Once again, communication theory is concerned with the statistical properties of the "channel" connecting r and s, because, for most engineering purposes, it is this channel whose characteristics must be exploited in designing effective coding strategies. The theory looks at a statistical quantity that is a certain weighted average of the conditional probabilities of all signals that can be transmitted from s to r. It does not concern itself with the individual events (the particular signals) except as a basis for computing the statistical functions that define the quantities of interest.

I skip over these matters rather lightly here, because it should be obvious that, insofar as communication theory deals with quantities that are statistical *averages* (sometimes called *entropy* to distinguish them from real information), it is *not* dealing with information as it is ordinarily understood. For information as it is ordinarily understood, and as it must figure in semantic and cognitive studies, is something associated with, and *only* with, individual events (signals, structures, conditions). It is only the particular signal (utterance, track, print, gesture, sequence of neural discharges) that has a content that can be given propositional expression (the content, message, or information carried by the signal). *This* is the relevant commodity in semantic and cognitive studies, and content – *what* information a signal carries – cannot be averaged. All one can do is average *how much* information is carried. There is no meaningful average for the information that my grandmother had a stroke and that my daughter is getting married. If we can say *how much* information theses messages represent, then we can speak about their average. But this tells us nothing about *what* information is being communicated. Hence, the quantities of interest in engineering – and, of course, some psychophysical contexts (Attneave 1959; Garner 1962; Miller 1953) – are not the quantities of interest to someone, like myself, concerned to develop an account of *what* information travels from source to receiver (object to receptor, receptor to brain, brain to brain) during communication.

Nevertheless, though communication theory has its attention elsewhere, it does, as Sayre (1965) and others have noted, highlight the relevant objective relations on which the communication of genuine information depends. For what this theory tells us is that the amount of information at r about s is a function of the *degree of lawful* (*nomic*) *dependence* between conditions at these two points. If two conditions are statistically independent (the way the ringing of *your* telephone is independent of the ringing of *mine*), then the one event carries no information about the other. When there is a lawful regularity between two events, statistical or otherwise, as there is between your dialing my number and my phone's ringing, then we can speak of one event's carrying information about the other. And, of course, this is the way we *do* speak. The ring *tells me* (informs me) that someone is calling my number, just as fingerprints carry information about the identity of the person who handled the gun, tracks in the snow about the animals in the woods, the honeybee's dance about the location of nectar, and light from a distant star about the chemical constitution of that body. Such events are pregnant with information, because they depend, in some lawfully regular way, on the conditions about which they are said to carry information.

If things are working properly, the ringing of my phone *tells* me that someone has dialed my number. It delivers this piece of information. It does *not* tell me that your phone is ringing, even if (coincidentally) your phone happens to be ringing at the same time. Even if A dials B's number whenever C dials D's number (so that D's phone rings *whenever* A dials B's number), we cannot say that the ringing of D's phone carries information about A's dialing activities – *not* if this "correlation" is a mere coincidence. We cannot say this, because the correlation, being (by hypothesis) completely fortuitous, does not affect the conditional *probability* of A's dialing B's number, given that D's phone is ringing. Of course, if we *know* about this (coincidental) correlation (though *how* one could know about its *persistence* is beyond me), we can predict one event from a knowledge of the other, but this doesn't change the fact that they are statistically independent. If I correctly describe your future by consulting tea leaves, this is not genuine communication *unless* the arrangement of tea leaves somehow depends on what you are going to do, in the way a barometer depends on meteorological conditions and, therefore, indirectly on the impending weather. To deny the existence of mental telepathy is not to deny the possibility of improbable cooccurrences (between what A thinks and what B thinks A is thinking); it is, rather, to deny that they are manifestations of *lawful* regularities.

Communication theory only makes sense if it makes sense to talk about the probability of certain specific conditions given certain specific signals. This is so because the quantities of interest to communication theory are statistical functions of these probabilities. It is this *presupposed* idea that I exploit to develop an account of a signal's content. These conditional probabilities determine how much, and indirectly *what*, information a particular signal carries about a remote source. One needs only to stipulate that the content of the signal, the information it carries, be expressed by a sentence describing the condition (at the source) on which the signal depends in some regular, lawful way. I express this theoretical definition of a signal's (structure's) informational content (Chapter 3, p. 65) in the following way:

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)

My gas gauge carries the information that I still have some gas left, if and only if the conditional probability of my having some gas left, given the reading on the gauge, is 1. For the same reason, the discharge of a photoreceptor carries the information that a photon has arrived (perhaps a photon of a certain wavelength), and the pattern of discharge of a cluster of ganglion cells carries the information that there is a sharp energy gradient (a line) in the optic array (Lindsay & Norman 1972; Rumelhart 1977). The following comments explain the main features of this definition.

1. There are, essentially, three reasons for insisting that the value of the conditional probability in this definition be 1 – nothing less. They are:

a. If a signal could carry the information that s was F while the conditional probability (of the latter, given the former) was less than 1 (.9 say), then the signal could carry the information that s was F (probability = .91), the information that s was G (probability = .91), but *not* the information that s was F *and* G (because the probability of their *joint* occurrence might be less than .9). I take this to be an unacceptable result.

b. I accept something I call the xerox principle: If C carries the information that B, and B's occurrence carries the information that A, then C carries the information that A. You don't *lose* information about the original (A) by perfectly reproduced copies (B of A and C of B). Without the transitivity this principle describes, the *flow* of information would be impossible. If we put the threshold of information at anything less than 1, though, the principle is violated. For (using the same numbers) the conditional probability of B, given C, could be .91, the conditional probability of A, given B, also .91, but the conditional probability of A, given C, less than .9. The noise (equivocation, degree of nomic *in*dependence, or nonlawful relation) between the end points of this communication channel is enough to break communication, even though every link in the chain passes along the information to its successor. Somehow the information fails to get through, despite the fact that it is nowhere lost.

c. Finally, there is no nonarbitrary place to put a threshold that will retain the intimate tie we all intuitively feel between knowledge and information. For, if information about s's being F can be obtained from a signal that makes the conditional probability of this situation only (say) .94, then information loses its cognitive punch. Think of a bag with 94 red balls and 6 white balls. If one is pulled at random (probability of red = .94), can you *know* (just from the fact that it was drawn from a bag with that composition of colored marbles) that it was red? Clearly not. Then why suppose you have the information that it is red?

The only reason I know for *not* setting the required probability this high is worries (basically skeptical in character) that there are no (or precious few) conditional probabilities of 1 – hence, that no information is ever communicated. I address these worries in Chapter 5. They raise issues (e.g., the idea of a "relevant alternative") that have received some attention in recent epistemology.

2. The definition captures the element that makes information (in contrast, say, to meaning) an important *epistemic* commodity. No structure can carry the information that s is F unless, in fact, s *is* F. False information, misinformation, and (grimace!) disinformation are not varieties of information – any more than a decoy duck is a kind of duck. A glance at the dictionary reveals that information is related to intelligence, news, instruction, and knowledge – things that have an important connection to *truth*. And so it should be with any theoretical approximation to this notion. Information *is* an important commodity: We buy it, sell it, torture people to get it, and erect booths to dispense it. It should not be confused with meaning, despite some people's willingness to speak of anything (true, false, or meaningless) stored on a magnetic disk as information.

3. Information, as defined above, is an objective commodity, the sort of thing that can be delivered to, processed by, and transmitted from instruments, gauges, computers, and neurons. It is something that can be *in* the optic array,[1] on the printed page, carried by a temporal configuration of electrical pulses, and stored on a magnetic disk, and it exists there *whether or not anyone appreciates this fact or knows how to extract it*. It is something that was in this world before we got here. It was, I submit, the raw material out of which minds were manufactured.

The parenthetical k occurring in the definition above (and explained below) relativizes information to what the receiver already knows (if anything) about the possibilities at the source, but this relativization does not undermine the essential objectivity of the commodity so relativized (MacKay 1969). We still have the flow of information (perhaps not so much) without conscious agents who know things, but without a lawfully regular universe (no matter how much knowledge we assign the occupants), no information is ever communicated.

4. A signal's informational content is not unique. There is, generally speaking, no *single* piece of information in a signal or structure. For anything that carries the information that s is a square, say, also carries the information that it is a rectangle, a parallelogram, *not* a circle, a circle *or* a square, and so on. If the acoustic pattern reaching my ears carries the information that the doorbell is ringing, and the ringing of the bell carries the information that the doorbell button is being pressed, then the acoustic pattern also carries the information that the doorbell button is being pressed (xerox principle). The one piece of information is *nested* in the other. This, once again, is as it should be. The linguistic meaning of an utterance may be

unique (distinguishable, for instance, from what it implies), but not the information carried by that utterance. Herman's statement that he won't come to my party means, simply, that he won't come to my party. It doesn't mean (certainly not in any linguistically relevant sense of "meaning") that he doesn't like me or that he can speak English, although his utterance may well carry these pieces of information.

5. The definition of a signal's informational content has been relativized to k, what the receiver (in the event that we are talking about a communication system in which the receiver – organism or computer – already has knowledge about the possible conditions existing at the source) already knows. This is a minor concession to the way we think and talk about information. The k is dischargeable by recursive applications of the definition. So, for instance, if I receive the information that your knight is *not* on KB-3 (by some signal), this carries the information that it *is* on KB-5, *if* I already know that the other possible positions to which your knight could have moved are already occupied by your pieces. To someone lacking such knowledge, the same signal does not carry this information (though it still carries the information that your knight is not on KB-3). The less we know, the more pregnant with information must be the signals we receive if we are to learn.

6. There is, finally, the important fact, already mentioned, that the informational content of a signal is a function of the *nomic* (or law-governed) relations it bears to other conditions. Unless these relations are what philosophers like to call "counterfactual supporting" relations (a symptom of a background, lawful regularity), the relations in question are not such as to support an assignment of informational content (Dretske 1977). The reason my thermometer carries information about the temperature of *my* room (the information *that* it is 72°F. in the room), but not about your room though both rooms are at the same temperature, is that (given its location) the registration of my thermometer is such that it *would not* read 72°F. *unless* my room was at this temperature. This isn't true of your room.

This fact helps explain an (otherwise puzzling) feature of information and, ultimately, of the cognitive attitudes that depend on it (belief, knowledge). For it is by virtue of this fact that a structure (some neural state, say) can carry the information that s (a distal object) is F (spherical) without carrying the information that s is G (plastic), even though (let us suppose) all spheres (in the relevant domain) are plastic. If the fact that all spheres are plastic is sheer accident, not underwritten by any lawful constraint, then the neural state might depend on s's being spherical without depending, in the same way, on its being plastic. Another way of expressing this fact (dear to the heart of philosophers) is to say that the informational content of a structure exhibits *intentional* properties. By saying that it exhibits intentional properties, I mean what philosophers typically mean by this technical term: that the informational content of a signal or structure (like the content of a belief, a desire, or knowledge) depends, not only on the reference (extension) of the terms used in its sentential expression, but on their *meaning* (intension). That is, in the sentential expression of a structure's informational content, one cannot substitute coreferring (i.e., referring to the same thing, coextensional) expressions without (possible) alteration in content. Just as a belief that this man is my cousin differs from a belief that he is Susan's husband, despite the fact that Susan's husband *is* my cousin (these expressions have the same reference), the information (as defined above) that he is my cousin differs from the information that he is Susan's husband. A signal can carry the one piece of information without carrying the other.

We have, then, an account of a signal's informational content that exhibits a degree of intentionality. We have, therefore, an account of information that exhibits some of the attributes we hope eventually to be able to explain in our account of our cognitive states. Perhaps, that is, one can know that s is F without knowing that s is G, despite the fact that all Fs are G, *because* knowledge requires information, and one *can* get the information that s is F without getting the information that it is G. If intentionality is "the mark of the mental," then we already have, in the physically objective notion of information defined above (even without k), the traces of mentality. And we have it in a form that voltmeters, thermometers, and radios have. What distinguishes us from these more pedestrian processors of information is not our occupation of intentional states, but the sophisticated way we process, encode, and utilize the information we receive. It is our *degree* of intentionality (see Part III).

Knowledge

Knowledge is defined (Chapter 4) as information-caused (or causally sustained) belief. The analysis is restricted to perceptual knowledge of contingent states of affairs (conditions having an informational measure of something greater than 0) of a *de re* form: seeing (hence, knowing) that this (the perceptual object) is blue, moving, a dog, or my grandmother.

This characterization of knowledge is a version of what has come to be called the "regularity analysis" of knowledge (Armstrong 1973; Dretske 1969; 1971). It is an attempt to get away from the philosopher's usual bag of tricks (justification, reasons, evidence, etc.) in order to give a more realistic picture of what perceptual knowledge is. One doesn't need reasons, evidence, or rational justification for one's belief that there is wine left in the bottle, if the bottle is sitting in good light directly in front of one. One can *see* that it is still half-full. And, rightly or wrongly, I wanted a characterization that would at least allow for the possibility that animals (a frog, rat, ape, or my dog) could know things without my having to suppose them capable of the more sophisticated intellectual operations involved in traditional analyses of knowledge.

What can it mean to speak of information as causing anything – let alone causing a belief? (The analysis of belief, the propositional attitude most often taken as the subjective component of knowledge, is postponed until Part III.) Assuming that belief is some kind of internal state with a content expressible as s is F, this is said to be caused by the information that s is F, if and only if those physical properties of the signal by virtue of which it carries this information are the ones that are causally efficacious in the production of the belief. So, for instance, not just any knock on the door tells you it is your friend. The (prearranged) signal is three quick knocks, followed by a pause, and then another three quick knocks. It is that particular signal, that particular temporal pattern, that constitutes the information-carrying property of the signal. The amplitude and pitch are irrelevant. When it is this pattern of knocks that causes you to believe that your friend has arrived, then (it is permissible to say that) the *information* that your friend has arrived causes you to believe he has arrived. The knocks might also frighten away a fly, cause the windows to rattle, and disturb the people upstairs. But what has these effects is not the information, because, presumably, the fly would have been frightened, the windows rattled, and the neighbors disturbed by *any* sequence of knocks (of roughly the same amplitude). Hence, the information is not the cause.

In most ordinary situations, there is no explanatory value in talking about the information (in an event) as the cause of something, because there is some easily identifiable physical (nonrelational) property of the event that can be designated as the cause. Why talk of the information (that your friend has arrived) as the cause, when it is clear enough that it is the particular temporal patterns of knocks (or acoustic vibrations) that was the effective agent?

The point of this definition is not to *deny* that there are physical properties of the signal (e.g., the temporal pattern of knocks in the above example) that cause the belief, but to say *which* of these properties must be responsible for the effect if the resultant belief is to qualify as knowledge.[2] If the belief that your friend has arrived is caused by the knock, but the pattern of knocks is irrelevant, then (assuming that someone else could be knocking at your door), though you are caused to believe it by the knock on the door, you do not *know* your friend has arrived. Those properties of the signal that carry the information (that your friend has arrived) are not the ones that are causally responsible for your belief.

The need to speak in this more abstract way – of information (rather than the physical event carrying this information) as the cause of something – becomes much more compelling as we turn to more complex information processing systems. For we then discover that there are an indefinitely large number of different sensory inputs, having no identifiable physical (nonrelational) property in common, that all have the same cognitive outcome. The only way we can capture the relevant causal regularities is by retreating to a more abstract characterization of the cause, a characterization in terms of its relational (informational) properties. We often do this sort of thing in our ordinary descriptions of what we see. Why did he stop? He could see that he was almost out of gas. We speak here of the information (that he was almost out of gas) that is contained in (carried by) the fuel gauge pointer and *not* the fuel gauge pointer itself (which, of course, is what we actually see), because it is a property of this pointer (its position, not its size or color) carrying this vital piece of information that is relevantly involved in the production of the belief. We, as it were, ignore the messenger bringing the information (the fuel gauge indicator) in order to focus on what information the messenger brings. We also ignore the infinite variety of optical inputs (all of varying size, shape, orientation, intensity) in order to focus on the information they carry. Often we have no choice. The only thing they have in common is the information they bear.[3]

A belief that *s* is *F* may not itself carry the information that *s* is *F* just because it is caused by this information (thereby qualifying as knowledge). A gullible person may believe almost anything you tell him – for example, that there are three elephants in your backyard. His beliefs may not, as a result, have any reliable relation to the facts (this is why we don't believe him when he tells us something). Yet this does not prevent him from knowing something he observes firsthand. When he *sees* the elephants in your backyard, he *knows* they are there, whatever other signal (lacking the relevant information) might have caused him to believe this. If the belief is caused by the appropriate information, it qualifies as knowledge whatever *else* may be capable of causing it.

This definition of knowledge accords, I think, with our ordinary, intuitive judgments about when someone knows something. You can't know that Jimmy is home by seeing him come through the door, if it could be his twin brother Johnny. Even if it is extremely unlikely to be Johnny (for Johnny rarely comes home this early in the afternoon), as long as this remains a relevant possibility, it prevents one from seeing (hence, knowing) *that* it is Jimmy (though one may be caused to *believe* it is Jimmy). The information that it is Jimmy is missing. The optical input is equivocal.

Furthermore, this account of knowledge neatly avoids some of the puzzles that intrigue philosophers (and bore everyone else to death). For example, Gettier-like difficulties (Gettier 1963) arise for any account of knowledge that makes knowledge a product of some justificatory relationship (having good evidence, excellent reasons, etc.) that *could* relate one to something false. For on all these accounts (unless special ad hoc devices are introduced to prevent it), one can be justified (in a way appropriate to knowledge) in believing something that is, in fact, false (hence, not know it); also know that *Q* (which happens to be true) is a logical consequence of what one believes, and come to believe *Q* as a result. On some perfectly natural assumptions, then, one is justified (in a way appropriate to knowledge) in believing the truth (*Q*). But one obviously doesn't *know* *Q* is true. This is a problem for justificational accounts. The problem is evaded in the information–theoretic model, because one can get into an appropriate justificational relationship to something false, but one cannot get into an appropriate informational relationship to something false.

Similarly, the so-called lottery paradox (Kyburg 1961; 1965) is disarmed. If one could know something without the information (as here defined), one should be able to know *before the drawing* that the 999,999 eventual losers in a (fair) lottery, for which a million tickets have been sold, are going to lose. For they all *are* going to lose, and one knows that the probability of each one's (not, of course, *all*) losing is negligibly less than 1. Hence, one is perfectly justified in believing (truly) that each one is going to lose. But, clearly, one cannot know this. The paradox is avoided by acknowledging what is already inherent in the information–theoretic analysis – that one cannot know one is going to lose in such a lottery no matter how many outstanding tickets there may be. And the reason one cannot is (barring a fixed drawing) the information that one is going to lose is absent. There remains a small, but nonetheless greater than 0, amount of equivocation for each outcome.

There are further, technical advantages to this analysis (discussed in Chapter 4), but many will consider these advantages purchased at too great a price. For the feeling will surely be that one never gets the required information. *Not* if information requires a conditional probability of 1. The stimuli are *always* equivocal to some degree. Most of us know about Ames's demonstrations, Brunswik's ecological and functional validities, and the fallibility of our own sensory systems. If knowledge requires information, and information requires 0 equivocation, then precious little, if anything, is ever known.

These concerns are addressed in Chapter 5, a chapter that will prove tedious to almost everyone but devoted epistemologists (i.e., those who take skepticism seriously). An example will have to suffice to summarize this discussion.

A perfectly reliable instrument (or one *as* reliable as modern technology can make it) has its output reliably correlated with its input. The position of a mobile pointer on a calibrated scale carries information about the magnitude of the quantity being measured. Communication theorists would (given certain tolerances) have no trouble in describing this as a noiseless channel. If we ask about the conditional probabilities, we note that these are determined by regarding certain parameters as fixed (or simply ignoring them). The spring *could* weaken, it *could* break, its coefficient of elasticity *could* fluctuate unpredictably. The electrical resistance of the leads (connecting the instrument to the apparatus on which measurements are being taken) *could* change. Error would be introduced if any of these possibilities was realized. And who is to say they are not *possibilities?* There *might* even be a prankster, a malevolent force, or a god who chooses to interfere. Should all these possibilities go into the reckoning in computing the noise, equivocation, and information conveyed? To do so, of course, would be to abandon communication theory altogether. For this theory requires for its application a system of fixed, stable, enduring conditions *within* which the degree of covariation in other conditions can be evaluated. If every logical possibility is deemed a possibility, then everything is noise. Nothing is communicated. In the same manner, if everything is deemed a *thing* for purposes of assessing the emptiness of containers (dust? molecules? radiation?), then no room, pocket, or refrigerator is ever empty. The framework of fixed, stable, enduring conditions within which one reckons the flow of information is what I call "channel conditions." Possible variations in these conditions are excluded. They are what epistemologists call "irrelevant alternatives" (Dretske 1970; Goldman 1976).

And so it is with our sensory systems. Certainly, in some sense of the word *could*, Herman, a perfectly normal adult, could be hallucinating the entire football game. There is no logical contradiction in this supposition; it is the same sense in which a voltmeter's spring *could* behave like silly putty. But this is not a sense of *could* that is relevant to cognitive studies or the determination of what information these systems are capable of transmitting. The probability of these things happening is set at 0. If they remain possibilities in some sense, they are not possibilities that affect the flow of information.

This discussion merely accentuates the way our talk of information *presupposes* a stable, regular world in which some things can be taken as fixed for the purpose of assessing the covariation in other things. There is here a certain arbitrary or pragmatic element (in what may be taken as permanent and stable enough to qualify as a channel condition), but this element (it is argued) is precisely what we find when we put our cognitive concepts under the same analytical microscope. It is not an objection to regarding the latter as fundamentally information-dependent notions.

Perception

Perception itself is often regarded as a cognitive activity: a form of recognizing, identifying, categorizing, distinguishing, and classifying the things around us (R. N. Haber 1969). But there is what philosophers (at least *this* philosopher) think of as an *extensional* and an *intensional* way of describing our perceptions (Dretske 1969). We see the duck (extensional: a concrete noun phrase occurs as object of the verb) and we recognize it (see it) as a duck – see *that* it is a duck (intensional: typically taking a factive nominal as complement of the verb). Too many people (both philosophers and psychologists) tend to think about perception *only* in the latter form, and in so doing they systematically ignore one of the most salient aspects of our mental life: the *experiences* we have when we see, hear, and taste things. The experience in question, the sort of thing that occurs in you when you see a duck (without necessarily recognizing it *as* a duck), the internal state without which (though you may be looking at the duck) you don't *see* the duck, is a stage in the processing of sensory information in which information about the duck is coded in what I call analog form, in preparation for its selective utilization by the cognitive centers (where the *belief* that it is a duck may be generated).

To describe what object you see is to describe what object you are getting information about; to describe what you recognize it as (see it to be) is to describe what information (about that object) you have succeeded in cognitively processing (e.g., that it is a duck). You can see a duck, get information *about* a duck, without getting, let alone cognitively processing, the information that it is a duck. Try looking at one in dim light at such a distance that you can barely see it. To confuse seeing a duck with recognizing it (either as a duck or as something else) is simply to confuse sentience with sapience.

Our experience of the world is rich in information in a way that our consequent beliefs (if any) are not. A normal child of two can *see* as well as I can (probably better). The child's experience of the world is (I rashly conjecture) as rich and as variegated as that of the most knowledgeable adult. What is lacking is a capacity to exploit these experiences in the generation of reliable beliefs (knowledge) about what the child sees. I, my daughter, and my dog can all see the daisy. I see it as a daisy. My daughter sees it simply as a flower. And who knows about my dog?

There are severe limits to our information-processing capabilities (Miller 1956), but most of these limitations affect our ability to cognitively process the information supplied in such profusion by our sensory systems (Rock 1975). More information *gets in* than we can manage to digest and get out (in some appropriate response). Glance around a crowded room, a library filled with books, or a garden ablaze with flowers. How much do you see? Is all the information embodied in the sensory representation (experience) given a cognitive form? No. You saw 28 people in a single brief glance (the room was well lit, all were in easy view, and none was occluded by other objects or people). Do you believe you saw 28 people? No. You didn't count and you saw them so briefly that you can only guess. That there were 28 people in the room is a piece of information that was contained *in* the sensory representation without receiving the kind of cognitive transformation (what I call digitalization) associated with conceptualization (belief). This homely example illustrates what is more convincingly demonstrated by masking experiments with brief visual displays (Averbach & Coriell 1961; Neisser 1967; Sperling 1960).

Although it is misleading to put it this way, our sensory experience encodes information in the way a photograph encodes information about the scene at which the camera is pointed. This is *not* to say that our sensory experience is pictorial (consists of sounds, sights, smells, etc.). I don't think there are daisy replicas inside the head, although I *do* think there is information about – and in *this* sense a representation of – daisies in the head. Nor do I mean to suggest (by the picture metaphor) that we are *aware of* (somehow perceive) these internal sensory representations. On the contrary, what we perceive (what we are aware *of*) are the things represented by these internal representations (not the representations themselves), the things *about which* they carry information (see section on "The Objects of Perception" in Chapter 6).

I see a red apple in a white bowl surrounded by a variety of other objects. I recognize it as an apple. I come to believe that it is an apple. The belief has a content that we express with the words, "That is an apple." The content of this belief does not represent the apple as red, as large, or as lying next to an orange. I may have (other) beliefs about these matters, but the belief in question abstracts from the concreteness of the sensory representation (icon, sensory information store, experience) in order to represent it simply as an apple. However, these additional pieces of information *are* contained in the sensory experience of the apple. As Haber and Hershenson (1973) put it (in commenting on a specific experimental setup), "It appears as if all of the information in the retinal projection is available in the iconic storage, since the perceiver can extract whichever part is asked for."

In passing from the sensory to the cognitive representation (from seeing the apple to realizing that it is an apple), there is a systematic stripping away of components of information (relating to size, color, orientation, surroundings), which makes the experience of the apple the phenomenally rich thing we know it to be, in order to feature *one* component of this information – the information that it is an apple. Digitalization (of, for example, the information that *s* is an apple) is a process whereby a piece of information is taken from a richer matrix of information in the sensory representation (where it is held in what I call "analog" form) and featured to the exclusion of all else. The difference between the analog and digital coding of information is illustrated by the way a picture of an apple (that carries the information that it is an apple) differs from a statement that it is an apple. Both represent it *as* an apple, but the one embeds this information in an informationally richer representation. Essential to this process of digitalization (the essence of conceptualization) is the *loss* of this excess information.

Digitalization is, of course, merely the information–theoretic version of stimulus generalization. Until information is deleted, nothing corresponding to recognition, classification, or identification has occurred. Nothing distinctively cognitive or conceptual has occurred. To design a pattern-recognition routine for a digital computer, for example, is to design a routine in which information *inessential* to *s*'s being an instance of the letter *A* (information about its specific size, orientation, color) is systematically discarded (treated as noise) in the production of some single type of internal structure, which, in turn, will produce some identificatory output label (Uhr 1973). If all the computer could do was pass along the information it received, it could not be credited with recognizing anything at all. It would not be responding to the essential sameness of different inputs. It would be merely a sophisticated transducer. Learning, the acquisition of concepts, is a process whereby we acquire the ability to extract, in this way, information from the sensory representation. Until that happens, we can see but we do not believe.

Belief

The content of a belief, what we believe when we believe (think) that something is so, can be either true or false. If we think of beliefs as internal representations (as I do), then these representations must be capable of *mis*representing how things stand. This is one aspect of intentionality.

Furthermore, if two sentences, S_1 and S_2, mean something different, then the belief we express with S_1 is different from the belief we express with S_2. Believing that a man is your brother is different from believing that he is my uncle (even if your brother is my uncle), because the sentences "He is your brother" and "He is my uncle" mean something different. A difference in meaning is sufficient, not necessary, for a difference in corresponding beliefs. The belief you express with the words "I am sick" is different from the belief I express with these words, despite the fact that the words mean the same thing. They have a different reference. This is a second aspect of intentionality.

But beliefs not only have a content exhibiting these peculiar intentional characteristics; they also, in association with desires, purposes, and fears, help to determine behavior. They are, if we can trust our ordinary ways of thinking, intentional entities with a hand on the steering wheel (Armstrong 1973).

It is the purpose of Part III to give a unified, information–theoretic account of these entities. The account is incomplete in a number of important ways, but the underlying purpose is to exhibit the way meanings (insofar as these are understood to be the conceptual contents of our internal states) are developed out of informational contents.

We have already seen (Chapter 3) the way information-bearing structures have a content (the information they carry – e.g., that *s* is *F*) exhibiting traces of intentionality. But this is only what I call the first order of intentionality. If two properties are lawfully related in the right way, then no signal can carry information about the one without carrying information about the other. No structure can have the (informational) content that *s* is *F* without having the (informational) content that *s* is *G*, if it turns out that nothing *can* be *F* without being *G*. This is the first respect in which the informational content of a structure fails to display the degree of intentionality of a belief (we can certainly believe that *s* is *F* without believing that *s* is *G*, despite the nomic connection between *F* and *G*).

The second respect in which information-carrying structures are ill prepared to serve as beliefs, despite their possession of content, is that, as we have seen, nothing can carry the information that *s* is *F*, nothing can have this informational content, unless, in fact, *s is F*. But we can certainly believe that something is so without its being so.

Without the details, the basic strategy in Part III is quite simple. Consider a map. What makes the symbols on a map *say* or *mean* one thing, not another? What makes a little patch of blue ink on a map mean that there is a body of water in a specific location (whether or not there actually *is* a body of water there)? It seems that it acquires this meaning, this content, by virtue of the information-carrying *role* that that symbol (in this case, a *conventionally* selected and used sign) plays in the production and use of maps. The symbol *means* this because that is the information it was designed to carry. In the case of maps, of course, the flow of information from map-maker to map-user is underwritten by the executive fidelity of the map-makers. A type of structure, in this case blue ink, means there is water there, even though particular instances of that (type of) structure may, through ignorance or inadvertence, fail to carry this information. Misrepresentation becomes possible, because instances (tokens) of a structure (type) that has been assigned (and in this sense has acquired) an information-carrying role may fail to perform in accordance with that role. The instances mean what they do by virtue of their being instances of a certain *type*, and the structure type gets its meaning from its (assigned) communicative function.

Neural structures, of course, are not conventionally assigned an information-carrying role. They are not, in this sense, symbols. Nevertheless, they acquire such a role, I submit, during their development in learning (concept acquisition). In teaching a child what a bird is, for example, in giving the child this concept (so that the youngster can subsequently have beliefs to the effect that this is a bird, that is not), we expose the child to positive and negative instances of the concept in question (with some kind of appropriate feedback) in order to develop a sensitivity to the kind of information (that *s* is a bird) that defines the concept. When the child can successfully identify birds, distinguish them from other animals (how this actually happens is, as far as I am concerned, a miracle), we have created something in the child's head that responds, in some consistent way, to the information that something is a bird. When the learning is successful, we have given the pupil a new concept, a new capacity, to exploit in subsequent classificatory and identificatory activities. If the child then sees an airplane and says "bird," this stimulus has triggered another token of a structure type that was developed to encode the information that the perceptual object was a bird (thereby representing it *as* a bird). We have a case of misrepresentation, a false belief.[4]

But we still have not captured the full intentionality of beliefs. In teaching our child the concept *water*, for instance, why say that the structure that develops to encode information about water is not, instead, a structure that was developed to encode information about the presence of oxygen atoms? After all, any incoming signal that carries the information that *s* is water carries (nested in it) the information that *s* has oxygen atoms in it (since there is a lawful regularity between something's being water and its having oxygen atoms in it).

The answer to this question is, of course, that the child has *not* developed a sensitivity to the information that *s* has oxygen atoms in it just because the pupil has been taught to respond positively to signals *all* of which carry that information. This can easily be demonstrated by testing the child with samples that are not water but do have oxygen atoms in them (rust, air, etc.). The crucial fact is that, although every signal to which the child is taught to respond positively carries information about the presence of oxygen atoms, it is not the properties of the signal carrying *this* information to which the child has acquired a sensitivity. Recall, it is those properties of the signal that are causally responsible for the child's positive response that define what information he is responding to and, hence, what concept he has acquired when he has completed his training. These properties (if the training was reasonably successful) are those carrying the information that the substance is water (or some approximation thereto – as time goes by, the concept may be refined, its information–response characteristics modified, into something more nearly resembling our mature concept of water).

Concept acquisition (of this elementary, ostensive sort) is essentially a process in which a system acquires the capacity to extract a piece of information from a variety of sensory representations *in* which it occurs. The child sees birds in a variety of colors, orientations, activities, and shapes. The sensory representations are infinitely variegated. To learn what a bird is is to learn to recode this analogically held information (that *s* is a bird) into a single form that can serve to determine a consistent, univocal response to these diverse stimuli. Until such structures have been developed, or unless we come into this world with them preformed (see the discussion of innate concepts in Chapter 9), nothing of cognitive significance has taken place.

NOTES

1. Though I am sympathetic to some of the (earlier) views of the late James Gibson (1950; 1966), and though some of my discourse on information (e.g., its availability in the proximal stimulus) is reminiscent of Gibson's language, this work was not intended as support for Gibson's views – certainly not the more extravagant claims (1979). If criticized for getting Gibson wrong, I will plead "no contest." I wasn't trying to get him right. If we disagree, so much the worse for one of us at least.

2. This is not so much a denial of Fodor's (1980) formality condition as it is an attempt to say *which* syntactical (formal) properties of the representations must figure in the computational processes if the resulting transformations are to mirror faithfully our ordinary ways of describing them in terms of their semantical relations.

3. I skip here a discussion of information's *causally sustaining* a belief. The idea is simply that one may already believe something when one receives the relevant supporting information. In this case, the belief is not caused or produced by the information. It nonetheless – after acquisition of the relevant information – qualifies as knowledge if it is, later, causally sustained by this information.

4. In my eagerness to emphasize the way conceptual content is determined by etiological factors (the information–response charactersitics of the internal structures) and to contrast it with the (behavioristically inspired) functionalist account (where *what* you believe is largely determined by the kind of output it produces), I seriously misrepresented (in Chapter 8) Dennett's (1969) position. Dennett stresses, as I do, the importance of the way these internal structures *mediate* input and output. He does, however, trace their ultimate significance, meaning, or content to the kind of (appropriate) behavior they produce.