

# The Metaphysics of Information

FRED DRETSKE, DURHAM, NC

## 1. “Information”

If you think information is important — and I assume most people do — you must have some vague idea of what it is. Since enormous amounts of money are spent on its collection, storage, and retrieval, since people are paid (and, during wartime tortured) to divulge it, somebody must know what it is. Somebody must know what all the fuss is about. It is easy enough to find people who *think* they know what it is, but very hard to find two people who agree. I was recently involved in a panel discussion the topic of which was: What Is Information?<sup>1</sup> There was a computer scientist, a mathematician, an electrical engineer, a biologist, a librarian, and psychologist, a linguist, and, yes, even a philosopher (me). No one had trouble talking for ten minutes. No two people said the same thing. They weren’t even close.

Maybe it isn’t important to know what information is. I have been told that this is a philosopher’s question (the typical “What is X?” question), the answer to which the rest of the world can ignore with impunity. Or, if people really do need an answer, it is quite acceptable for everyone — or perhaps each discipline — to have its own answer. The man on the street means this, the computer scientist that, the mathematician another thing, and the librarian still something else. If the philosopher wants to play this game, he can make up his own definition. Or he can try to figure out what everyone else means. All that is really necessary — and this is necessary only if people want to talk to and understand each other — is that people know what others mean.

I don’t think this will do. It leaves it a mystery why everyone thinks information is important or, if everyone who uses this word is really talking about *different* important things, why they insist on using the word “information” for *their* important thing. Why don’t they use some other word — like, maybe, “brick”, “fence post” or “vitamin”? Bricks, fence posts, and vitamins are important too, but nobody refers to them as information. Why not? Because we know what bricks, fence posts, and vitamins are, and they

clearly aren't information. Why is this so clear? Why can't a brick or a fence post be information? Because information is an epistemologically important commodity, and bricks and fence posts, though important, are not important in that way. Information has something to do with (perhaps it is even essential for) knowledge, and since knowledge is important, so is information. That is why people who would like a piece of the knowledge business — and that is a lot of people these days — want their product to be (or at least be intimately related to) information. So everybody ends up talking about his or her product as information.

This, I know, is a cynical view about the conceptual (at least verbal) mess in this area. The cynicism may or may not be justified. I don't know. But I do know there is a mess in this area and, as a result, a lot of confusion in the information sciences about what, exactly, information is and, therefore, what these sciences are supposed to be sciences of. With that in mind, what I hope to do in this paper is not to give a theory of information (I tried to do that elsewhere),<sup>2</sup> but to offer something more modest. I want to say what, at a commonsense level, information is. Wittgenstein (*On Certainty*, §260) said he preferred to reserve the expression "I know" for the cases in which it is used in normal linguistic exchange. I think that is good advice. I will try to say as clearly as I can what, in normal linguistic exchange, information is. If we are to build a theory of information, if there is ever to be a science of information, that, after all, is what we want a theory, a science, of — whatever we, in normal conversation, are talking about when we talk about information.

Information, as ordinarily understood, has three essential properties: (1) it is an intentional, a semantic, entity; (2) it is true; and (3) it is transmissible. Anything that lacks one or more of these properties isn't information.

## 2. Intentionality

Information is always information *about* something. If it isn't about anything, it isn't information. That is why bricks and fence posts aren't information. They aren't about anything. If someone or something sends or receives information, it always makes sense to ask, and there always has to be an answer to, the question, "Information about what?"

The only reason I can see why anyone would deny this is if they confused a signal with the information a signal carries — confused, that is,

information with the objects and events that carry it. A brick isn't information, but it might carry information. It might, for instance, carry information in its shape, color, size, or material composition, about its place and manner of manufacture. Or it might, by its deliberate placement and orientation, tell (inform) us where an accomplice is hiding. Similarly, we can use a fence post as a rough sundial — thus using the post (and its shadow) as a source of information about the time of day. Clearly, though, the information provided is not the brick, the post, or the shadow. It is what the brick, post, or shadow tell us about something else — in this case the time of day, place of manufacture, or location of an accomplice.

We have long been warned not to confuse words with what these words mean or refer to. The word “red” isn't the color red. Why, then, conflate the electrical charges in a silicon chip, a gesture (a wink or nod), acoustic vibrations, or the arrangement of ink on a newspaper page with what information these conditions convey? Syntax, organized facts about the bearers of meaning and information, is one thing; semantics, facts about the meaning or information they bear, is something altogether different. Just like meaning, information, is an abstract entity. It exhibits intentionality. It has aboutness. It comes in propositional form. Bearers of information, on the other hand, the objects and events (signals) that carry information, are, of necessity, concrete, observable entities. They are not about anything.

Information I give someone about my sister — that she lives in Colorado — isn't the same as information you give them about your sister — that she lives in Colorado — even if these pieces of information are embodied in physically indistinguishable forms — viz., the words “My sister lives in Colorado.” The reasons the information is different is not simply because the words come from different sources. Words can come from different sources and be the same. I can tell someone that my sister lives in Colorado and you can, by using different words, tell them the same thing (“Fred's sister lives in Colorado”). No, what makes the information I convey by telling someone my sister lives in Colorado different from what you convey by using these same words is not that the signals are different (they are, in fact, qualitatively the same), but that their reference is different. We are talking about different people. The intentionality is different. Different propositions are being expressed. That is why tree stumps can say something different — communicate different information — even when they are indistinguishable. The rings in this tree stump say that *this tree* is forty years old; the

rings in that other stump say that *that tree* is forty years old. A completely different piece of information.

### 3. Truth

Not only must information be about something, what it says about what it is about must be true for it to count as information. If it isn't true, it isn't information. Not every proposition is information. Only the true ones are (and not even all of them are — see § 4, Transmissibility).

I know we sometimes talk about misinformation and false information. This leads some people to conclude, mistakenly, that information needn't be true since, clearly, false information isn't true. This, I think, is a pretty heavy-handed treatment of ordinary language. It is like concluding that not all ducks are animals because decoy ducks aren't animals. The right conclusion to draw from the existence of decoy ducks, of course, is not that ducks don't have to be animals, but that decoy ducks aren't ducks. Likewise, the proper conclusion to draw from the existence of false information is not that information needn't be true. It is, rather, that false information is not information. False information is *fake* information and fake information is not a species of information any more than fake diamonds are a kind of diamond or phony dollar bills are (real) dollar bills.

I think that what fosters this tendency to think that information needn't be true is the mistaken assimilation of information to representation. Despite its falsity, a misrepresentation (a false representation) is nonetheless a representation. What is said is false, but something is still said. Things are verbally represented to be a certain way even when they are not the way they are represented to be. If you do not know whether Judith represented her whereabouts correctly, you still know she represented her whereabouts. You just don't know how well, how fully, how accurately, how truthfully, she represented it. Representation is like linguistic meaning. Words can mean that I'm at home even when I'm not at home. Words, and the people who utter those words, can represent things to be so that are not so. Misrepresentation is a species of representation. But information is not like this. Misinformation is not a piece of information of the false variety. It is not information at all. If you don't know whether what Judith told you about her whereabouts is true or not, you don't know whether she gave you information about her whereabouts. If she says she was at home on the

night of the crime, you know she *said* she was home. She *represented* herself as being home on that night. But you don't know — not yet anyway — whether her testimony provided information about her whereabouts that night. You won't know whether it provided information until you find out whether what she said is true — whether, in fact, she was at home that night.

If you doubt this, think about the role information plays in your daily life and what makes it such a valuable commodity. When you seek information in an airport or train terminal about how to get to Kirchberg am Wechsel from Vienna you aren't interested in merely being entertained by meaningful statements on the topic of Austrian geography. What you want are meaningful *true* statements. That is why you describe what you seek as information. You seek information because you want to know, and information provides what is needed to know: the truth. That is why in your efforts to find out how to get to Kirchberg you consult people who already know or can quickly find out. They already have the truth about these matters and so they can impart it to you. That is what information booths are for. If a person at the information booth tells you that Kirchberg am Wechsel is 400 miles due north of Vienna, you have a legitimate complaint. They cannot excuse their incompetence or deception by telling you that you only asked for information. You didn't ask for *true* information. That would be silly. In asking for information, one is asking for the truth.

This is part (but only a part) of the reason information is important. It is important because truth is important and information gives you the truth. That is why, from an epistemological standpoint, information is so much more important than meaning. Meaning can be either true or false. "Kirchberg is north of Vienna" is as meaningful as "Kirchberg is south of Vienna", but both statements can't be information because information has to be true, and both statements can't be true.<sup>3</sup>

I have said that information is what one needs to know, and that that is why it is important. It is, however, not *all* one needs to know. One may receive information about how to get to Kirchberg from Vienna and not be able to access (decode) it. The clerk at the information booth gives me the directions in German and I don't understand German. So although I've been given information about where Kirchberg is, I don't yet know where Kirchberg is. What I need besides the information is a good German-English dictionary. I need a way of extracting this information from the written or spoken message I have been given.

Having the information without being able to access it, without knowing what the information is or, perhaps, without even knowing that it is information, is a common enough occurrence. An expert on radioactivity tells me the half-life of radium. I don't realize he is an expert, and I don't believe him. I receive information and it doesn't generate knowledge because, in this case, it doesn't generate belief. Or I, a novice in the laboratory, observe the litmus paper turn blue and remain ignorant of the fact that the liquid in which it is placed is an acid. I don't realize that the paper carries this kind of information about acidity. Once again, the information is there in the condition I observe to be so, but I don't understand the language. I need a translation manual. In this case, I need to learn a little chemistry.

So information isn't sufficient for knowledge. Necessary, but not sufficient. It provides something, the truth, that is required for knowledge, but more is needed.

Skeptics, or simply advanced thinkers who don't like to talk about truth, will not like this result. They won't like it because, they will say, we cannot know what is true. If, then, information must be true to be information, we cannot know whether we have information or even whether there even is information to be had. Information, they will insist, is a useless commodity if it is impossible to know when you've got it.

This, I suspect, is the attitude of computer scientists who prefer to speak of anything that can be included in a database, whether it is true or false, as information. Computers, of course, can't distinguish true propositions from false propositions. Maybe (or so skeptics are inclined to think) our brains can't either. Feed a proposition in — that Kirchberg is 400 miles due north of Vienna, for instance — and the computer (or the brain), if it doesn't already have information that conflicts with this proposition, will treat it as information, a secure point for reasoning, problem solving, and inference. Everything that comes in, true or false, is treated as grist for the computational mill. For purposes of computation, for purposes of reaching a conclusion from the propositions existing in memory, the false is as good as the true. So, for computational purposes, for purposes of understanding the operation of computers (including the brain) why (one may ask) limit information to the true?

If your only interest is understanding computation, there is no reason to treat truth as an important commodity and, therefore, no reason to value information. Truth doesn't change anything. Valid arguments don't need

truth. False propositions work just as well. You can digitalize the false as easily as the true. But — and this is the important point — cognition in its most general sense is not just a matter of computation. It's not just a matter of determining that A follows from B, that C is probable relative to D, or that E is the best explanation of F. Computation is an instrument in a broader cognitive enterprise whose objective is getting things right, reaching true conclusions from available premises. That is why, for cognitive purposes, information is important. It provides what is needed in order for computation to get you where you want to go — the right answer, a true conclusion. Without information computation is worthless. We have a brain, yes, and it is fair to say that the brain's primary business is computation. So, for purposes of understanding the operation and function of the brain, just as for purposes of understanding the operation and function of silicon computers, truth — and thus information — isn't that important. It is dispensable.

Besides a brain, a computational organ, though, humans also have eyes and ears, and it is surely the business of these organs to supply the brain with what it needs — information — to reach *true* conclusions about the world. Without information, computational excellence is biologically worthless. What good is flawless reasoning if everything you conclude is false? That, indeed, is why information, though not particularly relevant to the business of computer (or, indeed, library) science, is absolutely essential to the cognitive life of the animals who use computers and libraries. Unlike machines and libraries, animals need to know. They need information. Computers and libraries don't. Nothing in the operation of a computer or of a library need change if we imagine everything passing through it to be false. It is different with animals.

## 4. Transmissibility

We send and receive information. We transmit it over telephone lines, by letters and photographs, in magazines and newspapers, and through the air. Light reflected from objects and sound emanating from them is full of information. That is why we have eyes and ears. If information couldn't move from one place to another, if it wasn't transmissible, it would be a useless commodity. It needs to be transferred to where it is needed — most often, in here, in the heads of living creatures, so it can be used in the service of need and desire satisfaction.

People (and dictionaries) sometimes confuse information with knowledge, but the transmissibility of information should be enough to distinguish these two items. I know things and you know things, and we can teach each other things. You can tell me what you know and give me knowledge. That is what education is about. We may — and sometimes do — speak of this as transferring or communicating knowledge, but we do not, not literally, transfer or transmit knowledge. It isn't knowledge that travels down the telephone line or (via electromagnetic waves) through the air. You don't give me your knowledge when you teach me something. You've still got it. You give me something that enables me to know, but what passes between us isn't knowledge. If it were knowledge, philosophers would be confused in their claims that belief was an ingredient in knowledge since what passes between teacher and student (a typed note, acoustic vibrations, light) clearly doesn't believe anything. So it isn't knowledge that is transmitted from one person to another. What is it then? It is information, of course: that which, by providing what is needed to know, gives knowledge to the person who receives it.

You may be wondering, though, how an abstraction like information can be moved through space? We can carry a brick or a fence post, good concrete space-time entities, from one place to another. We can transmit electrical current and electromagnetic waves from one place to another. But propositions? True propositions? How do you move a proposition, an abstract entity, from Chicago to Vienna? How do propositions, true propositions, entities that don't exist in space, change spatial location?

This may not seem like much of a problem. After all, we know how to transmit meaning, and meanings are as abstract, as non-spatial, as information. Write down what you mean, put it (what you've written) in a stamped envelope, deposit it in the mailbox, and — bingo! — meaning gets from Chicago to Vienna in a few short days. The meaning of the ink marks doesn't change in transit. So, in a sense, the meaning goes with the ink marks. It isn't lost. It may be thought, therefore, that the same is true of information. Information travels with the objects and events that carry information. Move these objects — written words, say — from Chicago to Vienna, and the information they carry (if there is any) goes with it. When you transmit a proposition, a meaning that is either true or false, you don't, when it is true, leave its truth behind. So you don't leave information behind.

Unfortunately, this quick answer won't do. The reason it won't is that information isn't *just* a true proposition. If that were all it was, there would be no problem transmitting it. You would do it the same way you do it with meaning. But information about X, although it is a true proposition about X, isn't *just* a true proposition about X. If something is information, it is true, yes, but there are true propositions about X that aren't information about X. To see why consider a familiar example in epistemology. A broken (12 hour display) clock tells the right time twice a day. Twice a day, that is, it says something true. But it *never* conveys information. You can never learn what time it is from such a clock even when it says something true. Broken clocks do not "tell" the time, not even twice a day. They do, however, twice a day, impart the truth.

I know next to nothing about quantum electrodynamics. So if you are looking for information about quantum electrodynamics, you are well advised to seek it elsewhere. Nonetheless, though I have no knowledge about this subject, I can quite easily communicate as many truths about it as the most knowledgeable expert. I merely take every proposition about this topic and assert both it and its negation. Half of what I say will be true. Unlike the expert, though, *none* of my declarations will convey information. You are not going to learn anything about electrodynamics from me. You will learn something from the expert and this despite the fact that I've told you everything he has. The communication of truth, though it is necessary for, is not sufficient for the communication of information. Something more is needed. What might this be?

In order to transmit information you must transmit not only the truth, but also what it takes for a suitably prepared receiver to *know* the truth you transmit. The relevant epistemic credentials, what it takes (besides truth) to produce knowledge, must also be communicated. That is why we use measuring instruments, gauges, indicators, alarms, and such — instruments that are reliably connected to the facts they provide information about. When working properly, these devices not only give one the truth about things one (typically) does not observe for oneself, but they connect one, by means of some reliable mechanism, to the facts in a way that enables one to learn what these facts are. You don't see the gas in your automobile gas tank, but the fuel gauge you do see, if it is working properly, carries information about how much gas there is in the tank. You can, therefore, come to know how much gas there is by consulting the gauge. The instrument creates a reliable connection, a channel, between you and the facts you seek informa-

tion about. This is why, in our quest for knowledge, we seek information from the experts, the people who know. People who know have the requisite credentials. They are, in this sense, our gauges and meters. They put us in touch with the facts we seek information about. Having observed things for themselves, or having obtained information from people who have, they are connected to the facts we seek information about in a way that makes their utterances authoritative. Their utterances are, in this way, like the pointer readings on a reliable indicator. That is why a genuine informant not only delivers the truth (an ignorant meddler can do this much). He or she creates a connection between source and receiver that enables the receiver to know the truth. Ignorant meddlers don't do that.

Information is transmitted, then, when this knowledge-conferring connection or channel between source and receiver, between facts to be known and a potential knower, is extended through space. Extending it is what copper wires, radio waves, light, sound, and newsprint allows us to do. I know my car keys are in my pocket because I can, by putting my hand in my pocket, feel them. In this case, tactile perception makes what happens in my brain depend on what is in my pocket. Information about my pocket is transmitted via nerves to my brain. I can, if I choose, now extend the causal influence of the keys in my pocket by telling you what is in my pocket. That makes what happens in your brain depend, to some extent, on what is in my brain and, therefore, on what is in my pocket. Information is transmitted. You now know where my keys are.

What is left for a genuine theory of information, it seems to me, is some systematic and revealing account of what this extendable, and knowledge conferring, relationship is. What, exactly, is it we must communicate with the truth to make the truth we communicate information? All I hope to have told you here is that this, indeed, is what we are looking for when we look for a theory of information. That is what a science of information should be a science of.

## References

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

Wittgenstein, L. 1969. *On Certainty*. Oxford: Blackwell.

## Notes

1. At Duke University, 2004.
2. Dretske 1981.
3. For more on the distinction between meaning and information, see Dretske 1981, pp. 41-44.

