

Information, Relevance, and Objectivity

Logan Fletcher

Philosophy and Cognitive Science
University of Toronto

Fred Dretske opens his book *Knowledge and the Flow of Information* with the assertion that information is *objective*. In contrast to the common view of information “as something that depends on the interpretive efforts—and, hence, on the prior existence—of intelligent life,” Dretske intends to develop a view of information “as an objective commodity, something whose generation, transmission, and reception do not require or in any way presuppose interpretive processes,” something “independent of its actual or potential use by some interpreter” (Dretske, 1981, p. vii).¹ Dretske has a significant stake in the question of information’s objectivity, because he is trying to use information to naturalize the mind:

Can you bake a mental cake using only physical yeast and flour? The argument is that you can. Given the sort of information described... something the most reflective materialists should be willing to give, we have all the ingredients necessary for understanding the nature and function of our cognitive attitudes. (1981, p. xi)

And, as Dretske points out elsewhere, “One cannot have a recipe for cake... that lists a cake as an ingredient.... This is why a recipe for thought cannot have interpretive attitudes or explanatory stances among the eligible ingredients” (2000, p. 209). In other words, if information is *not* objective, if it does depend crucially on interpretation, this effectively renders Dretske’s whole naturalization project viciously circular. As we might expect, then, Dretske takes up the challenge, in the remainder of the book and in other writings, of defending his claim that information is an “objective commodity”, by arguing that everything else on which information depends is itself objective.

¹ Throughout, ‘objective’ will be taken in Dretske’s sense to mean *interpretation-independent*, and will take on a more specific meaning as we progress. There may of course be other sorts of objectivity that *are* interpretation-dependent, but these will not be discussed.

It is surprising, then, to say the least, when in chapter 5, Dretske admits that whether or not a signal carries a piece of information is “a question that may not have an objectively correct answer... an issue that is, in part at least, responsive to the interests, purposes, and, yes, values of those with a stake in the communication process” (1981, p. 132-3). This would seem, on the face of things, to be utterly inconsistent with Dretske’s earlier claim that information is objective. Information, one might think, either *is* or *is not* objective. So what is going on here? Has this inconsistency simply gone unnoticed by Dretske? It has not; Dretske has a rather subtle story to tell about how these apparently opposing positions can be reconciled, and how we can consistently hold that information is indeed objective *in the sense required* for its use in a non-circular naturalization of the mind, while admitting that information is, *in a distinct sense*, relative to subjective interpretation. It is the topic of the present paper to examine this story, to see whether it holds up to close scrutiny. I will argue that it does not: The required separation between the subjective and objective aspects of information cannot be maintained; in the case of information, we cannot sequester the subjective factors in such a way as to maintain a space of objectivity purified of subjective influence. Subjective interpretation is inextricably entangled with every aspect of the informational phenomenon, leaving us no way to ‘factor out’ the subjectivity by relativizing information to an explicit standard. This has the result, I will argue, that Dretske’s project of using information to naturalize the mind is unavoidably circular.

Absolute Information

The idea of information, as understood by Dretske, is fairly easy to grasp. Dretskean information is roughly equivalent to Grice’s (1957) *natural meaning*, the sense in which *smoke means fire*. In Dretske’s terms, *smoke carries the information* that there is fire. Or, to take one of Dretske’s own examples, a pressure gauge carries information about the pressure within a certain boiler (1981, p. 123). In general, a *signal* (e.g. the pressure gauge being such-and-such) carries the information that conditions at the *source* are so-and-so (e.g. that the pressure is so-and-so) if and only if it is *certain*, given the signal being such-and-such, that the conditions at the source are indeed so-and-so. Here is how Dretske explicitly defines the informational content carried by a signal about a source:

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) (1981, p. 65)

Loewer points out that Dretske's invocation of probabilities is problematic, because there is no available account of probability that will meet Dretske's needs, and Dretske himself fails to provide one (Loewer, 1983). As Loewer acknowledges, however, the appeal to probability *per se* is not crucial to Dretske's account, and as Dretske himself explains, "A conditional probability of 1 between r and s is a way of describing a lawful (exceptionless) dependence between events of this sort," which "*nominally precludes* r 's occurrence when s is not F " (1981, p. 245). Dretske only ever distinguishes two "probabilities", namely $p=1$ and $0<p<1$, which, as far as I can tell, might just as well be called "necessity" and "possibility". What is important, then, is the nomic regularity that relates r to s , which need not be explained in terms of probabilities.

A more serious concern is the presence of " k " in the definition, which represents "what the receiver already knows... about the possibilities that exist at the source" (Dretske, 1981, p. 65). The motivation for this is the idea that if I already know that certain possibilities at the source are excluded, and you do not, I may receive a piece of information from the signal that you fail to receive from it. Dretske considers a shell game in which a peanut is under one of four shells: If I know the peanut is not under shells 1 and 2 (I have seen they are empty), and you do not know this, then, when we both see that shell 3 is empty, I receive the information that the peanut is under shell 4, whereas you do not (1981, p. 78-9). The definition of informational content in terms of knowledge would seem to compromise the supposed objectivity of information; in fact, it appears to lead straightforwardly to circularity, since Dretske goes on to analyze knowledge as information-caused belief (1981, p. 86). Dretske anticipates this potential criticism, arguing that the apparent circularity is in fact a finite regress, in which we can recursively analyze any presupposed knowledge into its informational basis, until we eventually "reach the point where the information received does not depend on any prior knowledge about the source" (1981, p. 86-7). At the point where this recursive process bottoms out, we will be left with what Steven Savitt calls "absolute informational content", which is defined just as in Dretske's original definition, but without the parenthetical reference to k (Savitt, 1987). Cohen and Meskin (2006) express doubts about whether this recursive process will always bottom out into absolute information in the way Dretske suggests, but it is worth noting that there is a much more direct way to reduce knowledge-dependent information to absolute information, which is simply to consider the absolute *correlate* of the knowledge-relative informational content: the informational content carried by the same signal *without* taking k into

account. To be sure, this is not a proper ‘reduction’, since the absolute correlate will generally be less specific in its content than the knowledge-dependent information, but, as Savitt notes, we can still explain cases such as the discrepancy between observers in the shell game, by positing an inferential process by which we derive the knowledge-dependent content (that the peanut is under shell 4) from the absolute content (that shell 3 is empty) plus our prior knowledge (that shells 1 and 2 are empty). There are, then, good reasons to doubt that the inclusion of *k* in Dretske’s definition seriously compromises the objectivity of information, so defined. But in any case, Dretske’s theory includes a theory of absolute information. Therefore, I will set aside the question of *k*-relativity for the rest of the paper, restricting my attention exclusively to *absolute information*. Should *this* fail to be objective, it will create much more serious problems for Dretske’s project.

Channel Conditions

As Dretske’s definition makes clear, for a signal to carry information about a source *just is* for the appropriate nomic regularity to hold between them: *Given* that the signal is such-and-such, *necessarily* the source is so-and-so. Crucially, this nomic regularity supports the following counterfactual claim: *Were it not the case that* the source is so-and-so, *it would not be the case that* the signal is such-and-such. It is this lawful quality—indicated by ability to support counterfactuals—that distinguishes informational relationships from mere *de facto* correlations, a distinction stressed by Dretske as crucial (1981, p. 73-5). Dretske at one point glosses these nomic regularities as “natural laws” (1981, p. 77-8), but they are strikingly different from familiar laws of nature. Laws of physics, for instance, are thought to express necessary relationships *because* they have never been observed to admit of exceptions. In contrast, it is strange to think that there is something *necessary* about the relationship between the pressure gauge and the boiler pressure. We do our best to maintain this relationship, of course, but it is not difficult to imagine cases in which it fails to hold, for instance when the gauge is damaged, or disrupted by a strong magnetic field, or simply disconnected—that these possibilities are *not* the case is of course not something we can tell simply by reading the gauge itself. It would therefore seem to be a contingent matter whether or not the signal tracks conditions at the source, not any sort of law.

Dretske explains that the informational relationship between signal and source holds only relative to a set of *channel conditions* (1981, p. 115). In the example of the pressure gauge, the set of channel conditions will consist of all the facts, other than the fact that the boiler pressure is so-and-

so, on which the such-and-such-ness of the pressure gauge depends: That the gauge is hooked up properly to the boiler, that there are no strong magnetic fields interfering, and so forth. It would seem that the complete set of channel conditions would have to be infinite, since no matter how long our list grows, there will always be *some* scenario, however far-fetched, in which the signal is such-and-such, all the channel conditions already listed are met, and *still* the source is not so-and-so. Now, one might expect Dretske to say that if each of these facts is *true* (that the gauge is hooked up properly, that there are no strong magnetic fields, etc.), and the signal is such-and-such, *then* the signal will carry the information that the source is so-and-so. Dretske does not say this, and for good reason. The problem is not that the source will not be so-and-so (it will), but that, in the envisioned scenario, there is no way to distinguish the source from the channel conditions. Since changing *any* one of these conditions will affect the signal, there is no objective fact about which one is actually the source—we can treat any of them as the source, depending on our interests.² Clearly, this sort of relativism will be unacceptable to Dretske.

Here is what Dretske does say: The channel conditions (for absolute information) are the “existing conditions (on which the signal depends) that... generate no (relevant) information” (1981, p. 115). What does it mean for a condition to “generate no (relevant) information”? It means, first, that the fact associated with the condition (“there are no magnetic fields interfering”) is *actually* true, and, second, more problematically, that there are no *relevant alternative possibilities* to it being true. For a channel condition to *hold*, then, requires more than just something being true of the actual world—there is an additional *modal* requirement. Hence it not only must be *true* that there is no magnetic field interfering, but also the *possibility* of a magnetic field interfering must not be relevant. It is by appealing to this modal aspect that Dretske is able to distinguish between the source and the channel conditions; the source has relevant alternative possibilities (which nomically covary with conditions at the signal), whereas the channel conditions do not.

What does it mean for a possibility to be relevant or irrelevant? If there is a magnetic field that unpredictably turns ON and OFF, interfering with the gauge’s readings only when it is ON, then Dretske would surely say that the gauge will not *ever* carry information about the boiler pressure, *even*

² We cannot, however, treat *all* of them as the source at the same time, since there must be a channel in order for there to be a source (Dretske, 1981, p. 117). It makes no sense for the signal to carry information about all the conditions on which it depends.

when the magnetic field is *OFF*, because it will *always* be a relevant possibility that it is *ON*. On the other hand, if a magnetic field has *never* interfered with the boiler gauge, then Dretske would probably consider the possibility of magnetic interference to be an irrelevant one. As Dretske acknowledges (1981, p. 130), however, *de facto* frequencies of this sort, covering actual past occurrences, are not criterial of whether or not something is a relevant possibility. Indeed, they *cannot* be criterial, since this would undermine Dretske's central claim that informational relationships are *nomic* regularities, as opposed to *de facto* correlations. So what *does* make something a "relevant possibility"?

Relational Objectivity

Whether or not the signal carries information about the source depends on whether or not the channel conditions hold. Hence, it must be an objective truth that the conditions hold, if the information-carrying relationship they support is itself to be objective. But as we've seen, Dretske is forced to include a modal element in the definition of channel conditions: Whether or not a channel condition *holds* is in turn dependent upon what relevant alternative possibilities are present in the actual situation. Does it make sense for a *possibility* to be *actually* present, and for this to be an objective fact? Dretske thinks it does:

The difference between a relevant and an irrelevant alternative resides, not in what we happen to *regard* as a real possibility (whether reasonably or not), but in the kind of possibilities that actually exist in the objective situation. (Dretske, 1981a, p. 377)

Dretske discusses some examples that he takes to illustrate the actual presence and absence of relevant possibilities. For instance, Dretske considers a bird-watcher who identifies a bird as a Gadwell duck. Supposing that the bird actually *is* a Gadwell, does the bird-watcher *know* that (i.e. possess the information that) it is a Gadwell?³ Is the possibility that it is a look-alike Siberian Grebe a *relevant* alternative? (1981a, p. 369) It depends, Dretske argues, on the actual facts: The possibility is surely relevant if there are in fact Grebes in the vicinity. On the other hand, it is irrelevant, according to Dretske, if Grebes do not actually exist, or if "due to certain geographical barriers, they are confined to their Siberian habitat" (1981a, p. 377). Then there are borderline cases, which we can produce *ad*

³ This is, of course, a question about k-relative, not absolute, information. This is incidental, however—the point being made concerns relevant alternative possibilities in general.

nauseum: What if the Grebes remain in Siberia, but there is no barrier *preventing* them from migrating? (1981a, p. 377) What if a hunting organization has plans to populate the area with Grebes, but has not yet been given governmental permission to do so? How serious must these plans be, to give rise to a relevant alternative? Suppose they have already moved caged Grebes into the area—do we trust them to keep the birds locked away safely? What if one of the handlers is known to be careless? And what *sort* of carelessness, exactly, is appropriate to generalize to apply to the specific case of Grebe-handling?

It is cases such as these, apparently unbounded in their potential for subtle contextual entanglements, that impel Dretske to acknowledge a pragmatic element to the presence of the information-carrying relationship. Dretske admits that such cases make information-carrying “a question of degree” (1981, p. 132), but it is difficult not to feel that in admitting this, he is admitting too little. Dretske describes “the difficult question” as “the question of when an alternative... is just *too remote* to qualify as relevant” (1981a, p. 376), but it would seem that the more challenging pragmatic problem would be to determine the metric of remoteness that this formulation presupposes. Is a Grebe migration from Siberia *more remote* a possibility than a hunting organization realizing their half-hearted plan to import Grebes? Whichever way your intuitions may tend on this matter, it is undoubtedly a question whose answering requires the balancing of two content-rich complexes of reliability-cum-variability that are not commensurable with each other in any *straightforward* way. In contrast, the pragmatic problem discussed by Dretske—that of determining what *degree* of remoteness is sufficient to render a possibility irrelevant—would seem to be precisely analogous to the problem of determining *how* tall a person must be in order to be considered a “tall” person, a question that may well involve some interesting sensitivity to context (relativizing to age, gender, or race, for instance), but which clearly lacks the rich entanglements of questions regarding the relative remoteness of possibilities.

Bearing this in mind, here is how Dretske resolves the issue of the acknowledged pragmatic element to relevant possibilities: He claims that *being a relevant alternative possibility* (hence also *being an information-carrying relationship*) is an absolute matter, determined relative to a pragmatically-constituted standard of *what counts as relevant*. The case is analogous, Dretske argues, to the question of whether something is *flat*: Although ‘flat’ is an “absolute concept”, meaning that “nothing can be flat if it has *any* bumps and irregularities,” flatness is determined relative to “what *counts* as a bump or irregularity” (1981a, p. 366). Dretske calls concepts of

this sort “relationally absolute”, and claims that relevant possibilities and information-carrying relationships have a relationally absolute character (1981a, p. 367). Information-supporting nomic regularities, then, are indeed exceptionless, permitting no relevant possibilities, but they must be understood relative to a standard of relevance.

So much for the absoluteness of information; what about its objectivity? Is it relational in the same manner? Dretske’s position here is not so easy to discern. On the one hand, when faced with the difficulties of determining the relevant possibilities in challenging borderline cases, Dretske admits that the “question may not have an objectively correct answer” (1981, p. 132-3), thereby suggesting that what objectivity it does have will need to be understood relationally. And when Dretske speaks of seeking relevant alternatives in the “possibilities that actually exist in the objective situation” (1981a, p. 377), the “objective situation” is clearly one constituted by a standard of relevance. Otherwise, relevant possibilities would exist *independent* of such a standard, and then what would we need the standard for? On the other hand, Dretske states that a skeptic, “someone with very high standards, someone who considers almost any alternative relevant,” would “speak falsely” by making claims about (the failure of) information transmission on the basis of such high standards (1991, p. 192). But if relevant possibilities are determined *relative to a standard of relevance*, then on what basis does Dretske assert that the skeptic’s standard is mistaken, if not by appealing to a second-order standard that renders the first-order standard redundant? Dretske’s answer is that when one makes claims about the success or failure of information transmission, it is the “background of standard or intended *use* [of information] to which the claim is relative... A speaker who brings a novel or unusual viewpoint to the situation being described... is expected to indicate that difference in order to avoid misunderstanding” (1991, p. 194).⁴ Dretske, then, is not providing an argument that the skeptic is *objectively wrong* to choose an eccentric standard of relevance; such an argument would be absurd, since the appeal to “standard or intended *use*” as an *objective* means of disqualifying the skeptic’s standard would directly contradict Dretske’s description of information, *qua* objective, as “independent of its actual or potential use by some interpreter” (1981, p. vii). Rather, Dretske is making the thoroughly

⁴ David Lewis makes the same point in a very similar context, arguing that a man who “attached eccentric relative importances to respects of comparison of worlds” would not be entitled to make claims on the basis of this standard, “at least not without giving warning of his eccentric notions.” Lewis claims the eccentric individual “lies”, because “he temporarily changes the conventional meaning of his words” (Lewis, 1979, p. 93-4).

pragmatic claim that the skeptic's standard lies outside the range of conventional standards, which are more *reasonable* for identifying information *as it is conventionally used*.

The objectivity of information, then, is indeed to be understood relationally, as relative to a standard of relevance, just as the absoluteness of information is so relative. The skeptic's determinations of relevant possibilities, then, will be objective *relative to* the skeptic's standard, just as any other set of determinations will be objective *relative to* its own standard. This means that the standard favored by Dretske will necessarily be arbitrary from the perspective of objectivity itself—this is an unavoidable consequence of understanding relevant possibilities in relation to a standard of what *counts* as relevant. Nevertheless, Dretske can still argue, and I take him to be arguing, that the standard that is pragmatically reasonable (or, perhaps, any standard within the set of pragmatically reasonable standards) provides a basis for objective claims about the transmission of information. Just because the standard of relevance that effectively defines information transmission is *chosen* on the basis of pragmatic issues regarding the use of information, does not imply that information, so defined, is not objective. Were the definition *itself* to make essential reference to information use, objectivity would surely be compromised, but this is a completely different matter from information use playing an essential role in *motivating the choice* of definition. The difference is that between defining length in terms of “reasonable units” and defining length in terms of “metres” or “feet”—the latter standards, unlike the former one, provide a foundation for objective claims about length, in spite of the subjective and pragmatic concerns that may have figured in the decision to adopt them in the first place.

As suggested by the example of length, a paradigmatically objective property, *all* objective claims will at some level exhibit this standard-relativity. For a claim to be objective is just for there to be a clean interface—an unambiguous standard—joining together the subjective matter of *what counts* for deciding the claim, and the objective matter of how the claim is decided, *given what counts*, in a manner that enforces their mutual separation. Subjectivity comes into play in the *formulation* of a standard, but the *application* of a standard, once formulated, needs no further help from subjective interpretation—it applies, so to speak, all on its own. For example, in playing a game, we can have diverging subjective opinions regarding *what should count* as a legal or illegal move.⁵ Such subjective

⁵ Conossieurs of spontaneous game invention will be familiar with cases in which the agreed-upon rules of a newly invented game turn out not to anticipate every situation that

matters, however, are (in most games, at least) strictly limited to issues concerning the formulation of the rules, which, once formulated, provide a clear standard that objectively determines issues of legality. There is then a sharp distinction between the subjective, value-laden, and, I am tempted to say, *aesthetic* issues involved in determining *what game we are playing*, and the purely objective issues about what is and is not the case with respect to legality, *given that we are playing this game*. The rules, once formulated, can be applied *objectively*, that is to say, without any concern for the subjective issues that figure in why the rules are as they are. There is a clear line where subjectivity leaves off, and objectivity takes over.

This, I think, is how Dretske wants us to understand the nature of information as simultaneously pragmatic and objective. Dretske wants to say that subjective issues determine the standard by which possibilities are judged to be relevant (what game we are playing), and that the facts about facts about which ones are relevant (which moves are legal), *given this standard* (given that we are playing this game), are entirely objective, leaving no leeway for subjective interpretation. But we might reasonably worry that this sort of “objectivity” is vacuous, since we can *always* posit a standard according to which our claims are “objective”, regardless of how paradigmatically *subjective* these claims are. Consider, for instance, if Dretske had argued that *beauty* is objective, *relative to a standard of what counts as beautiful*. Our own interests and values, he says, figure in the determination of our personal standard of beauty, but *given this standard*, there are perfectly objective facts about what is and is not beautiful. I take it that Dretske does not want to argue that information is objective *in the same sense that beauty is*, but what prevents us, then, from saying that beauty is relationally objective? The answer, I think, is that *there is no such thing* as a relationally objective standard of beauty, not even of Dretske-beauty, beauty as judged by Dretske. Beauty is so richly intertwined with subjectivity, so unavoidably tacit, that it is just not possible to ‘factor out’ the subjective element by formulating a standard that can be applied in abstraction from this subjective vantage point. Anything that could genuinely be said to distinguish beautiful things would have to *appreciate* this beauty, and hence would need to be, in some sense, a subject, with an *ineliminably implicit* grasp of aesthetic value. Or so, at any rate, I would claim.

Whether or not you agree with me about beauty, I think it is clear that *if* the evaluation of a claim is indeed dependent on such ineliminably

emerges within play. There may in such cases be opposing sets of strong feelings regarding which potential extension of the rules is most “true to the spirit” of the game.

implicit elements, the claim cannot be said to be an objective one. This consideration leads directly to what we might call the *Explicitness Condition*: In order for a standard to serve as a basis for making (relationally) objective claims, that standard must be able to be formulated in an explicit manner. There must be no implicit factors determining *how* the standard is to be applied; such implicitness must be cashed out, by rendering these factors explicitly, as part of the formulation of the standard. If the application of the standard depends on factors that are systematically resistant to explicit formulation, then this suggests that the standard fails as a basis for making objective claims. The idea behind the Explicitness Condition is that if a standard is to support relational objectivity, its application can neither depend upon nor leave room for subjective interpretation. If it is in fact interpretation-independent, it ought to be amenable to explicit expression.⁶

The question of the objectivity of Dretskean information, then, comes down to the question of whether the conventional or pragmatically reasonable intuitions about relevant alternative possibilities can be captured in a standard that meets the Explicitness Condition. Dretske himself does not attempt to articulate anything approaching such a standard, and no available theories of relevant alternatives seem to make any progress towards discharging the implicit and pragmatic aspects. David Lewis's theory of relevant alternatives (Lewis, 1996), for instance, described as "the most sophisticated version... to date" (Shaffer, 2001, p. 202), makes essential reference to "salient resemblance" in analyzing relevant alternatives, which does not leave us any better off with respect to the task of eliminating the dependence on pragmatic elements. It is worth pointing out that for most theorists, Dretske included, the primary use to which relevant alternatives are put is in arguments against the skeptical claim that knowledge is nonexistent or severely limited. As Stewart Cohen points out,

⁶ There will of course be skeptical worries about whether *any* standard could ever be formulated in an entirely explicit manner. The skeptic will claim that there will always be factors on which the application of the standard depends, that remain necessarily implicit. To some extent, this concern will be addressed by the discussion of the mechanization of the application of standards in the last section of this paper. For now, I will just point out that we can, if we wish, interpret the Explicitness Condition as a regulative ideal, so that the Condition just requires that the standard can be rendered *arbitrarily* explicit, depending on how carefully and conscientiously we proceed in its formulation. On this interpretation, even if any actual formulation of a standard will always depend on residual implicit elements, the standard can still meet the Explicitness Condition.

the notion of a relevant alternative can do *this* job perfectly well, even if we are unable “to provide a precise formulation of the criteria of relevance” (1991, p. 33), because here there is no prior commitment to the interpretation-independent character of knowledge (or information). If, on the other hand, “we want the theory of relevant alternatives to provide an analysis of knowledge”—as Dretske of course does, via the notion of information—“then the failure to provide precise criteria would constitute a failure of the theory” (Cohen, 1991, p. 33).

Although it seems to me that the burden lies with Dretske to show that we *can* formulate an explicit standard for determining relevant alternatives, the lack of any extant accounts that accomplish this surely does not demonstrate that we *cannot*. There are, however, reasons to doubt that we can formulate such a standard, reasons to think that the presence of relevant alternatives, like the presence of beauty, is an inexorably subjective matter. For one thing, in spite of Dretske’s insistence that the issue of which possibilities are relevant is “a question of degree”, there is not any straightforward manner of assessing the degree of remoteness of possibilities. As we observed earlier, relative remoteness cannot be assessed without becoming embroiled in complex issues regarding the forms of reliability and variability that are present in situations. Dretske tells us that

To qualify as a relevant possibility, one that actually affects the equivocation of (and therefore the information in) a signal, the possibility envisaged must actually be realizable in the nuts and bolts of the particular system in question. (1981, p. 131).

The suggestive metaphor of “the nuts and bolts of the system” brings to mind the picture of certain basic pieces that, by virtue of their shapes, place constraints on the manner with which they can be combined. The problem is that in the real-world situations in which Dretske wants information to be carried, it is not obvious which is the relevant decomposition into nuts and bolts. What *counts* as a “bolt”, or as a “hinge point” is itself extremely sensitive to context. Recall the question about what *sort* of carelessness is appropriate to generalize to the case of Grebe-handling. It is difficult to see how such a question could be answered without delving into issues about the nature of carelessness. But it is also not clear that the *sort* of carelessness that is relevant to Grebe-handling is a detachable “piece” that would arise in other contexts as well. The nearness of the possibility that the Grebe-handler will fail to properly secure the Grebe’s cage may well depend on an idiosyncratic constellation of factors, including, for instance, competence with locks and cages. But what *sort* of cage-competence generalizes to the current situation? For *any* particular situation, we will be able to ask

indefinitely many such questions, regarding the *proper fit* of regularities to the concrete situation. The standard of relevance, though, if it is to serve as a basis for talking about “information” *in general*, must provide an answer, not only to all such questions in any particular situation, but to all such questions in *all* of the situations in which we want to talk about information. With similar considerations in mind, Cohen says that “Because there are so many complex and controversial variations in examples like this, it is exceedingly difficult to capture the distinction between relevant and irrelevant alternatives in a precise criterion” (Cohen, 1991, p. 32). Indeed, the task of formulating an explicit standard of relevance seems prohibitively difficult.

Modal Complexity

At this point, I want to suggest that the difficulties we encounter in evaluating claims about relevant possibilities can fruitfully be viewed as endemic to a larger class of claims, what I will call *modally complex* claims, of which relevant possibility claims are a subset. Setting aside all claims of an *obviously* normative or evaluative nature, such as moral or aesthetic claims, I propose that we can divide much of what remains into three classes. First, there are *non-modal* claims: These are claims about the *actual* world, stripped of all modal character, that is, independent of any facts about what *could be*, *would have been*, or *almost is* the case. We do not need to appeal to possible worlds at all in assessing non-modal claims. Claims about the mass of something, or the relative positions of two things, are paradigmatic examples. Second, there are *modally simple* claims, claims about what is *possible* or *necessary*, either absolutely or conditionally. Modally simple claims make no reference to the actual world, and hence can be assessed without considering what is *actually* the case. Absolute possibility, under most interpretations of modality, would be identified with *logical* possibility, since logical truths are usually thought to hold in all possible worlds. Claims about conditional possibility, for instance claims about what is *physically* possible, are directed to certain subsets of the set of possible worlds. Third, there are *modally complex* claims: These are claims directed to what we might think of as the *modally-animated actual world*, that is, the actual world, considered as situated amongst surrounding possibilities. Counterfactual conditionals and claims about the nearness and remoteness of possibilities are modally complex—they are true or false *at*

the actual world, but they can only be assessed by considering the relation of the actual world to other possible worlds.⁷

Non-modal claims and modally simple claims are both potentially capable of relational objectivity, that is, we can formulate standards, on the basis of which we can understand such claims, which meet the Explicitness Condition. For instance, claims about the lengths of various things can be objective relative to a standard that expresses how the terms by which we measure lengths are to be applied. Or claims about what is possible given, say, Newtonian physics can be understood relative to an explicit statement of the appropriate laws and fundamental entities (here we actually specify what the “nuts and bolts” are). In fact, the categories of non-modal and modally simple claims appear to exhaust the claims to which empirical science and logic, respectively, are committed.

The standards for evaluating modally complex claims, on the other hand, are systematically resistant to explicit formulation. This is a consequence of the fact that modally complex claims depend crucially on relations of similarity between the actual world and other possible worlds, which relations may shift dramatically depending on the particular claim at issue. David Lewis defends the view that the “relative importances of respects of comparison,” and hence also the relations of “comparative overall similarity” between worlds, while volatile, are “at least roughly fixed,” and can be treated as such for the purposes of evaluating modally complex claims (Lewis, 1973, p. 92-3). But consider the counterfactual conditional with which Lewis opens his book: “If kangaroos had no tails, they would topple over” (1973, p. 1). This claim reveals the basic nature of counterfactuals: A change is made to the actual world, which will have further consequences, creating a rupture that must be patched up in such a way as to preserve *relevant global similarity* to the greatest possible degree. But which is the best way to patch things up? Are we really to believe that a world in which toppling kangaroos miraculously survive the crunch of natural selection is *more similar* to ours than a world in which their bodies compensate in some way to regain their balance?⁸ Yet prior to raising such

⁷ Strictly speaking, modal complexity does not require any reference to the actual world, since we can treat other possible worlds as modally animated, and evaluate counterfactuals, for instance, at these worlds. But in the vast majority of cases of concern, the world at which modally complex claims are true or false is the actual world.

⁸ Snapping turtles, with their reduced shell on their underside (unlike painted turtles, which have full shells), have compensated by evolving a sharp beak, claws, and an aggressive disposition. This would seem to cast doubt on the counterfactual claim that “if painted turtles had a reduced lower shell, they would be rendered vulnerable”.

considerations, the intuitively obvious response is indeed that the counterfactual is true—they *would* topple over. In contrast, saying “if humans had no baby fingers, they would spurt blood from their hands” seems much less plausible, though it is not easy to say what makes the difference, much less to formalize the difference in an explicit standard. The lesson here is that the similarity relations between worlds, on which all assessments of all modally complex claims depend, are intimately sensitive to the particular claim being made, in ways that are subtle and difficult to discern.

Dretskean information, since it is stipulated to be distinct from mere *de facto* correlations, is inherently a modally complex phenomenon. For this reason, variations of Dretske’s theory, even if they avoid the problems that arise in dealing with channel conditions, will inevitably face difficulties of importantly the same sort. The “objective counterfactual theory” of Meskin and Cohen (2006), for instance, which takes the counterfactual *were the source not so-and-so, the signal would not be such-and-such* as definitive of the informational relationship, will succumb to all of the difficulties involved in evaluating counterfactuals.⁹ To see that these will be *the same* problems that confront Dretske, consider that Dretske’s channel conditions, since they support (at least certain cases of) this same counterfactual, can be understood as a roundabout way of providing a theory of its truth conditions. The modal complexity of information belies Dretske’s portrayal of information as “something the most reflective materialists should be willing to give,” the “physical yeast and flour” out of which we can build the mind (1981, p. xi)—physics, after all, is non-modal, and information is modally complex.

Intelligence and Information

Modal complexity, though confounding, is a natural and pervasive aspect of the human psychological world. The extent to which we experience the world as modally animated can sometimes make it difficult for us to discern the modal complexity in what we naïvely take to be the actual world. For instance, take Dretske’s comment that if there are Grebes “in the vicinity”, we do not know that we are seeing a Gadwell (1981a, p. 369). This is putatively a description of the *actual* world, on the basis of

⁹ Although Meskin and Cohen address some formal issues concerning counterfactuals, they do not discuss the more substantive concerns about the standards by which counterfactuals are evaluated, beyond asserting that they see “no serious reason to think the notion of counterfactuals is incoherent” (2006, p. 347). Coherence, needless to say, is not tantamount to objectivity.

which we can supposedly judge which alternative possibilities are relevant. But in fact “in the vicinity” is already a modally loaded notion; if someone is “in the area”, this just means, on a psychological level, that we *might* run into them. It is our naïve tendency to take the modally animated world of experience for the *actual* world that lends more plausibly to the idea of “the kind of possibilities that actually exist in the objective situation” (Dretske, 1981a, p. 377) than it deserves. If we think of actuality in terms of notions such as “in the vicinity”, as opposed to, say, “3010 metres away, north-by-north-east”, we will more readily accept the idea that the *possibility* that *this is a Grebe right here* is to be found “in the objective situation”.

In fact, the human ability to effortlessly perceive the world as modally animated might reasonably be taken to be a central and defining feature of intelligence itself. Douglas Hofstadter, for instance, has argued that “the crux of creativity resides in the ability to manufacture variations on a theme,” including “counterfactual conditionals”, “subjunctives”, and “‘almost’-situations” (i.e., near possibilities) (1985, p. 249). Consider the frame problem in artificial intelligence, characterized by Jerry Fodor as “to all intents and purposes... the problem of how the cognitive mind works” (1987, p. 148). The frame problem is the problem of predicting the relevant side effects of an event, while ignoring the irrelevant ones (see the discussion in Haugeland, 1985, p. 203-11). The problem is essentially that of finding the *relevant* true subjunctive conditionals that concern the immediate future, i.e., conditionals of the form *if X happened now, Y would be the case*. Although such conditionals are considered *logically* to be about the actual world (albeit in the future), the manner in which the future animates the present is *psychologically* similar to the manner in which nearby possible worlds animate the present in the cases of counterfactuals and, of course, relevant alternative possibilities. Zeroing in on relevant alternative possibilities, and bringing them to bear appropriately on the present situation, is something we humans are very good at—it plays a crucial role in our own assessments of the deliverances of perception. The problem of determining which possibilities are relevant, which, as we have seen, requires the *appropriate* application of past regularities to the idiosyncratic demands of the present situation, as well as the identification of the regularity and variability in the particular situation, the nuts and bolts and hinge points we perceive to be present in our modally animated world—this problem is at least as deserving as the frame problem of being considered central to intelligence. If there is *any* behavioral competence that is a sure sign of intelligence, it seems, to me at least, that the ability to discriminate relevant alternative possibilities is just as reasonable a

candidate as anything else we might think up. Like the frame problem, the problem of identifying relevant possibilities runs, in Fodor's memorable phrase, "as deep as the analysis of rationality" (1987, p. 140).

Dretske endorses a view that he states in the form of a slogan: "If you can't make one, you don't understand how it works" (Dretske, 1994). This is an expression of Dretske's view that the best test of a naturalistic theory of a mechanism is whether or not the theory tells us how to *construct* the mechanism. This view is put forward in much the same spirit as the Explicitness Condition, but it takes things one step further, telling us what *counts* as explicitness—if the theory of the mechanism is sufficiently explicit, enabling us to understand how it works, it should, *ipso facto*, tell us how we can build one. One of the virtues of Dretske's buildability condition is that, like the Explicitness Condition, it serves the function of removing implicit, and (hence) potentially subjectivity-concealing elements: If our theory of how a mechanism works depends tacitly on unthematized subjective factors, these will surely be forced to the surface if we attempt to build the mechanism on the basis of the theory.

Could we, then, use Dretske's buildability constraint to evaluate whether or not a standard serves as a suitable foundation for relational objectivity? This seems reasonable; if the standard is really explicit—that is, if it really provides a basis for making objective claims—we can, on the basis of the standard's explicit formulation, construct a mechanism that serves the function of *applying* the standard. If we really understand how the rules of a game work, we ought to be able to build a mechanism that indicates whether or not the rules are being followed. What happens, though, when we try to build a mechanism that applies the standard of relevance that is the basis for the supposed objectivity of relevant possibilities, and, ultimately, of information? Of course, this is highly implausible, since, as we've observed, explicit standards of relevance are not so much as a gleam in the eye of any theorist today. But suppose we *were* able to construct a mechanism that could tell whether or not a possibility was relevant. If the preceding considerations are correct—those regarding the centrality of the recognition of relevant alternative possibilities to the nature of intelligence—then the mechanism we would end up building would *be intelligent*. The standard of relevance, naturally resistant to explicit formulation, will end up embodied implicitly in the thinking of a *subject*, one with interests and values, who pays attention to relevant possibilities, while ignoring irrelevant ones. The result for Dretske's project is devastating, because he has got things exactly backwards: It is *information* that depends on *minds*, not the other way around.

Works Cited

- Cohen, J. and Meskin, A. 2006. An objective counterfactual theory of information. *Australian Journal of Philosophy*. 84: 333-52.
- Cohen, S. 1991. Skepticism, relevance, and relativity. In McLaughlin, B., Ed., *Dretske and his critics*.
- Dretske, F. 1981. *Knowledge and the flow of information*. Oxford: Basil Blackwell.
- Dretske. 1981a. The pragmatic dimension of knowledge. *Philosophical Studies*. 40: 363-78.
- Dretske. 1991. Knowledge: Sanford and Cohen. In McLaughlin, B., Ed., *Dretske and his critics*.
- Dretske, F. 1994. If you can't make one, you don't know how it works. In 2000, *Perception, knowledge and belief*. Cambridge, UK: Cambridge University Press.
- Fodor, J. 1987. Modules, frames, fridgeons, sleeping dogs and the music of the spheres. In Pylyshyn, Z., Ed., *The robot's dilemma: The frame problem in artificial intelligence*. Norwood, NJ: Ablex Publishing Co.
- Grice, P. 1957. Meaning. *Philosophical Review*. 66: 377-88.
- Haugeland, J. 1985. *Artificial intelligence: the very idea*. Cambridge, MA: MIT Press.
- Hofstadter, D. 1985. *Metamagical themas*. NY: Basic Books.
- Lewis, D. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Lewis, D. 1996. Elusive knowledge. *Australian Journal of Philosophy*. 74: 529-67.
- Loewer, B. 1983. Information and belief. *Behavioral and Brain Sciences*. 6: 75-6.
- McLaughlin, B., Ed. 1991. *Dretske and his critics*. Cambridge, MA: Basil Blackwell.
- Savitt, S. 1987. Absolute informational content. *Synthese*. 70: 185-90.
- Shaffer, J. 2001. Knowledge, relevant alternatives and missed clues. *Analysis*. 61: 202-8.