



Counterfactuals vs. conditional probabilities: A critical analysis of the counterfactual theory of information

Hilmi Demir

To cite this article: Hilmi Demir (2008) Counterfactuals vs. conditional probabilities: A critical analysis of the counterfactual theory of information, *Australasian Journal of Philosophy*, 86:1, 45-60, DOI: [10.1080/00048400701846541](https://doi.org/10.1080/00048400701846541)

To link to this article: <http://dx.doi.org/10.1080/00048400701846541>



Published online: 28 May 2008.



Submit your article to this journal [↗](#)



Article views: 109



View related articles [↗](#)



Citing articles: 5 View citing articles [↗](#)

COUNTERFACTUALS VS. CONDITIONAL PROBABILITIES: A CRITICAL ANALYSIS OF THE COUNTERFACTUAL THEORY OF INFORMATION

Hilmi Demir

Cohen and Meskin [2006] recently offered a counterfactual theory of information to replace the standard probabilistic theory of information. They claim that the counterfactual theory fares better than the standard account on three grounds: first, it provides a better framework for explaining information flow properties; second, it requires a less expensive ontology; and third, because it does not refer to doxastic states of the information-receiving organism, it provides an objective basis. In this paper, I show that none of these is really an advantage. Moreover, the counterfactual theory fails to satisfy one of the basic properties of information flow, namely the Conjunction principle. Thus, I conclude, there is no reason to give up the standard probabilistic theory for the counterfactual theory of information.

Introduction

Since the publication of Shannon's seminal article [1948], philosophers have used information theoretic concepts to explain such notions as mental content, belief, and knowledge. This approach reached a high point in the 1980s, with Dretske's *Knowledge and the Flow of Information* [1981], where he attempted to explain perceptual content, belief, and knowledge in terms of informational content. However, like many other things in life, the peak also signalled the beginning of the end. In the following twenty years, many have given up on this project, for two reasons. First, Dretske's account appeals to inverse conditional probabilities, a type of probability not supported by the standard interpretations of probability. Moreover, Dretske assigns unity to this inverse conditional probability, which requires him to deny the possibility of partial information and misinformation. Many found this constraint too demanding. As a result, information-theoretic concepts fell out of favour in the 1990s. Recent years, however, have seen a renewed interest in information-theoretic concepts, as many philosophers have offered new definitions of Dretske's problematic concepts of entropy and informational content [Usher 2001; Eliasmith 2000; Scarantino 2005; Cohen and Meskin 2006]. Here I focus on the most recent attempt: Cohen and Meskin's counterfactual theory of information [2006].

In his review of Dretske's book, Loewer [1983] surveys three standard interpretations of probability and argues that the type of conditional probability that Dretske wants is not supported by any of them. This problem, which I shall call the probability interpretation problem, presents a big challenge for information theoretic approaches. Cohen and Meskin address this difficulty by defining informational content in terms of counterfactuals, thus avoiding any reference to conditional probabilities. Then, they claim that their approach not only overcomes the probability interpretation difficulty but also brings about three important advantages.¹ The counterfactual definition, they say,

- i. Provides a better framework for information flow properties.
- ii. Requires a less expensive ontology.
- iii. Provides an objective basis because it does not refer to doxastic states of the information-receiving organism.

In this paper, I argue that none of these alleged advantages is really an advantage. Moreover, the counterfactual approach fails to satisfy some basic intuitions about information flow. Consequently, the only advantage their view has over Dretske's is avoiding the probability interpretation problem. But this is not a sufficient reason to prefer Cohen and Meskin's proposal to Dretske's. There are other ways to overcome the probability interpretation difficulty, possibilities that have not, in my view, been adequately explored in the literature. Although I cannot discuss these possibilities in detail here, in the last section I sketch two promising paths for this purpose. My main aim in this paper is to show that Cohen and Meskin are mistaken about the three advantages of their counterfactual based definition.

Before I proceed with my arguments, a terminological preamble is in order.

I. A Terminological Preamble

In order to use information-theoretic concepts to solve philosophical problems, one must first determine the informational content of a signal. A mental representation can be considered as a signal, and its informational content can be used to identify that mental representation's proper content. Likewise, a belief or a knowledge statement can also be considered as an information carrying signal. Thus the generalized question is how to define the informational content of a signal. Since Shannon's original article [1948], the trend has been to define informational content probabilistically. This can be done in two ways. Let's assume that a signal r is produced (or

¹My original understanding was that Cohen and Meskin advocated the counterfactual theory of information. However, they have stated via personal correspondence that their efforts were meant as exploration, not advocacy. Either way, their counterfactual theory of information is a valuable contribution to the literature. However, if their intentions are not to defend the counterfactual approach, then there is a slight mischaracterization of their intentions in my article. The reader is strongly encouraged to read their original paper [2006], and judge for himself.

caused) by an event s . One can define the informational content of r either as the conditional probability of r given s , or as the conditional probability of s given r . The first option defines the content in terms of the probability of the signal given the event that produces it; by contrast, the second defines the content in terms of the probability of the original event given the signal. In other words, the first uses the probability of the response given the stimulus, whereas the second uses the probability of the stimulus given the response. Table 1 depicts these ways of stating the difference between the first and the second options.

Table 1

$\text{Pr}(r s)$	$\text{Pr}(s r)$
Probability of the response given the stimulus	Probability of the stimulus given the response
Probability of the signal given the triggering event	Probability of the triggering event given the signal

The type of conditional probability depicted in the left hand column is known as ‘forward conditional probability’, because it follows the temporal order of the causal sequence. The stimulus which produces the response precedes it in time, and so the conditional probability of the response given the stimulus is going forward. By contrast, the type represented in the right hand column, the probability of the stimulus given the response, is called ‘inverse conditional probability’, because it proceeds from the ‘effect-response’ to the ‘cause-stimulus’. As applied to mental states, these two forms can be characterized as the probability of a mental state given the triggering event, as opposed to the probability of the external event given a mental state. These descriptions are depicted in Table 2.

Table 2

$\text{Pr}(r s)$	$\text{Pr}(s r)$
Forward Conditional Probability Probability of a mental state given the triggering external event	Inverse Conditional Probability Probability of the triggering event given a mental state.

In his theory of informational content, Dretske runs into the probability interpretation problem precisely because he uses inverse conditional probabilities. To explain this in detail is the task of the following section.

II. Dretske’s Probabilistic Definition and the Probability Interpretation Problem

Dretske [1981] defines the informational content of a signal, and uses this definition to explain representational notions like ‘seeing that’, ‘believing that’, and ‘knowing that’. His probabilistic definition is as follows.

Informational Content: A signal r carries the information that s is F if and only if the conditional probability of s 's being F , given r (and k), is 1 (but, given k alone, less than 1) [k refers to background knowledge].

[Dretske 1981: 65]

Because Dretske uses the probability of the event given the message, or $\Pr(s|r)$, his definition clearly opts for the inverse conditional interpretation.

In his review of Dretske's book, Loewer [1983] raises an important worry about Dretske's use of inverse conditional probabilities. He argues that none of the available interpretations of probability supports using inverse conditional probabilities for defining informational content. The subjective interpretation of probability as degree of belief does not work for Dretske since he wants an objective and naturalistic account of mental content and knowledge. Of the objective notions, relative frequency is also not suitable because it does not apply to single unrepeatable events, which Dretske also wants to include in his account of informational content. Although propensity interpretations do not run into the problem of singleton events, they support only forward conditional probabilities, as Loewer explains:

The propensity of a chance setup C producing outcome $[r]$ is usually explained... as a measure of the causal tendency of C producing $[r]$. But Dretske is after the converse probability, the probability that r was produced by a chance setup C . This probability is usually not meaningful on a propensity (or for that matter a frequency) interpretation. The point is that $P(r|C)$ may be meaningful but not $P(C|r)$, since there may be no propensity $P(C)$.

[Loewer 1983: 75 (my corrections in square brackets)]

This difficulty for Dretske's inverse conditional probability interpretation is what I am calling the probability interpretation problem.² Following Loewer, Cohen and Meskin cite this problem as their main motivation for offering a counterfactual definition of informational content. Let us now turn to their definition and the advantages they claim for it over Dretske's probabilistic definition.

III. Cohen and Meskin's Counterfactual Theory of Information

In their paper, Cohen and Meskin [2006] argue against using inverse conditional probabilities to define informational content, and offer an alternative definition based on counterfactuals. They claim that their definition has three advantages over the former. Their plan of attack, following Loewer's strategy in 1983, has two main steps.

Step 1: Show that the standard accounts of information in circulation use inverse conditional probabilities.

²In his book, Dretske [1981: 245] argues against the relative frequency interpretation. However, in his response to Loewer, he seems to be defending the relative frequency approach [Dretske 1983: 84]. To my knowledge, this tension in his work has not been pointed out before.

Step 2: Show that it is difficult to make sense of inverse conditional probabilities on any of the standard interpretations of probability.

Because of its influence, they focus on Dretske's theory as the paradigmatic standard account for their discussion. Cohen and Meskin cite Loewer's arguments that none of the available interpretations of probability grounds the inverse conditional probabilities Dretske's theory needs. They intend to generalize their results to other standard accounts of information, and in particular to Shannon's mathematical theory of communication. They claim that Shannon's mathematical theory of communication also runs into the probability interpretation problem, since Shannon's fundamental notion of mutual information makes 'ineliminable' reference to inverse conditional probabilities. Here is what they say:

The remarks that follow are applicable to other accounts of information (both semantic and quantitative) that are grounded in conditional probabilities. Most saliently, consider the setup of Shannon [1948]: let $\{s_1, \dots, s_n\}$ be discrete alternative states of a source s with probabilities $\{P(s_1), \dots, P(s_n)\}$ respectively, and let $\{r_1, \dots, r_k\}$ be discrete alternative states of a receiver r with probabilities $\{P(r_1), \dots, P(r_k)\}$ respectively; assume that $P(s_i) > 0$ for all i , that $P(r_j) > 0$ for all j , and that

$$\sum_{i=1}^n P(s_i) = \sum_{j=1}^k P(r_j) = 1.$$

Shannon defines *the mutual information between s and r* as follows:

$$I(s, r) = - \sum_{i=1}^n P(s_i) \log_2 P(s_i) + \sum_{j=1}^k P(s_j | r_j) \log_2 P(s_j | r_j).$$

So defined, mutual information makes ineliminable reference to the same sorts of inverse conditional probabilities as Dretske's theory, and so is vulnerable to the concerns we raise about the interpretation of those probabilities.

[Cohen and Meskin 2006: 336 n. 7]

The expression ' $P(s|r)$ ' in the quote refers to the probability of the stimulus given the response, i.e., the probability of an external state of affairs given a mental representation, which is clearly an instance of 'inverse conditional probability'. Therefore, Cohen and Meskin reason, the probability interpretation problem applies to Shannon's mathematical theory of communication as well.

What they miss, however, is that mutual information is commutative. That is, $I(s, r)$ is equal to $I(r, s)$. Hence, the formula for $I(s, r)$, the mutual information between s and r , can be rewritten as the following.

$$I(s, r) = I(r, s) = - \sum_{j=1}^k P(r_j) \log_2 P(r_j) + \sum_{i=1}^n P(r_j | s_i) \log_2 P(r_j | s_i).$$

In this formula, there is no reference to inverse conditional probabilities. The probabilities are all forward conditional probabilities which, as Cohen and Meskin would agree, can be given the propensity interpretation with no difficulty. Although it is true that Shannon himself favoured the relative frequency approach, there is no basis in his theory for rejecting the propensity interpretation. The point here is not which interpretation of probability should be used, but rather that Cohen and Meskin's claim about Shannon's theory is mistaken. Thus whereas Dretske's account runs into the probability interpretation problem, the same is not true of Shannon's mathematical theory of communication. This mistake does not affect Cohen and Meskin's main point, however, but only the scope of their criticisms. They could have just focused on Dretske's theory without applying their criticisms to other accounts. Their counterfactual approach makes a valuable contribution to the literature and deserves to be discussed on its own terms.

Cohen and Meskin begin with a crude version of their counterfactual theory, and then revise it by adding a non-vacuousness clause to avoid some difficulties concerning necessary truths. For both the crude and revised accounts, they present one weak and one strong version. The weak versions take the counterfactual criterion as only a sufficient condition for information-carrying relations whereas the strong versions take it as both necessary and sufficient. This difference between their strong and weak versions is irrelevant to their claim that the counterfactual account is preferable to Dretske's account. Hence, for the sake of simplicity I shall discuss only the weak version of their position.

Here is the weak version of their revised (non-vacuous) version of informational content:

(W*) x 's being F carries information about y 's being G if the counterfactual conditional 'if y were not G , then x would not have been F ' is non-vacuously true.

[Cohen and Meskin 2006: 335]

This 'non-vacuousness' clause excludes assigning the information-carrying relation to cases where y 's being G is necessarily true. If y 's being G is necessarily true, then the counterfactual will come out true no matter what, hence the counterfactual will be vacuously true. Following the generally accepted intuition that necessary truths carry no information at all,³ Cohen and Meskin aim to exclude necessary truths from the set of information-carrying signals by adding the non-vacuousness clause.

Cohen and Meskin argue that their counterfactual theory of information is preferable to Dretske's definition based on inverse conditional probabilities for three reasons. First, according to Dretske, information flow must be transitive; i.e., if A has the information B and B has the information C , then A has to have the information C . This 'intuitive' requirement leads to some unacceptable consequences, which the counterfactual theory avoids

³It is important to note that some philosophers disagree with this claim about necessary truths. Carnap and Bar-Hillel [1952], Hintikka [1970], and Bremer [2003] are useful sources for a balanced presentation of this debate.

by not requiring transitivity of information flow. I shall discuss this question in the following section. Second, they claim that Dretske's account makes essential reference to nomic regularities, whereas the counterfactual account is agnostic about whether information-carrying relations must appeal to nomic regularities. Thus, they conclude, the counterfactual definition requires a more economical ontology. The third advantage they claim for their counterfactual definition concerns the subjectivity of doxastic states. Because Dretske's theory makes essential reference to background knowledge for information-carrying relations, Cohen and Meskin claim that it cannot provide an objective, reductive explanation of mental representation. By contrast, their account makes no such reference to background knowledge and consequently it can provide an objective explanation of mental content. I claim that none of these reasons provides a substantial advantage for their view over Dretske's. Let us examine their arguments one by one.

A. Information Flow Properties

The most controversial feature of Dretske's definition [stated on p. 47] is assigning unity to the conditional probability. This leads to several unacceptable consequences, such as denying the possibility of partial information and misinformation. Despite these consequences, Dretske claims that he is obliged to assign unity, because it is the only way to match our common sense intuitions about information flow. Two of these intuitions are what he calls the Conjunction principle and the Xerox principle. According to the former, if a signal r carries the information that A and if it carries the information that B , then it has to carry the information that A and B . The latter is the transitivity property of information flow: if A has the information that B and B has the information that C , then A has to have the information that C . These claims are intuitively true, and any technical definition of the information-carrying relation must conform to them. On the other hand, we know from probability theory that conditional probabilities satisfy these two principles only when their values are 1.

Although Cohen and Meskin do not discuss the Conjunction principle, they claim that the Xerox principle, the transitivity property mentioned above, is true for most but not all cases. More precisely, they claim that the information-carrying relation is neither transitive nor intransitive but non-transitive. In fact in the literature it is pointed out that the Xerox principle holds only for informational chains which form Markov chains, and so it has a limited application [Demir 2006]. Cohen and Meskin believe that their account matches the limitation on the Xerox principle better than does Dretske's. In Dretske's framework the following argument is valid.

A has the information that B.

B has the information that C.

Therefore, A has the information that C.

However, the following inference schema is not valid in Lewis's [1973] possible worlds semantics which is commonly accepted as the standard account.⁴

$A \Box \rightarrow B$

$B \Box \rightarrow C$

Therefore, $A \Box \rightarrow C$

This is because the closest possible A-world may not be a C-world given that the closest possible A-world is a B-world and the closest possible B-world is a C-world. So even if the conclusion follows from the premises in many cases, there could be cases in which it does not. It is true that, as Cohen and Meskin claim, this result conforms better to the limitation on the Xerox principle. If this were the only result of the counterfactual account for the intuitive properties of information-carrying relation, then it would constitute a reason to prefer the counterfactual account over Dretske's. However, that is not the case. As we shall see, the conjunction principle mentioned above provides insurmountable difficulties for the counterfactual account.

There are good reasons for questioning the general application of the Xerox principle, but it is very difficult to come up with a reason to reject the Conjunction Principle. The Conjunction principle implies that if a signal r carries the information B and if it also carries the information C, then it has to carry the information B and C. It is one of Dretske's reasons for assigning unity to conditional probabilities. Cohen and Meskin do not discuss the Conjunction principle. In fact, their counterfactual definition fails to satisfy it. Let us assume that A carries the information that B and A also carries the information that C. According to their definition, A carries the information that B if the counterfactual 'if B were not the case, then A would not have been the case' is non-vacuously true. When this definition is applied to two assumptions, one gets the following counterfactual claims:

1. 'If B were not the case, then A would not have been the case' is true.
2. 'If C were not the case, then A would not have been the case' is true.

Now, the question is whether these two necessarily imply the following: 'If B were not the case and C were not the case, then A would not have been the case' is true. According to Lewis's possible world semantics, which is the canonical account according to Cohen and Meskin, the truth conditions for the two counterfactual claims, are the following:

1. Truth Condition of 1: the closest not-B-world is also a not-A-world.
2. Truth Condition of 2: the closest not-C-world is also a not-A-world.

⁴Here, I use Lewis' possible worlds semantics for two reasons: first, it is the standard account; and second, Cohen and Meskin [2006: 347] also accept it as the canonical account. By no means do I claim that it is the only appropriate semantics for counterfactuals.

These two truth conditions do not imply that the closest not-B and not-C-world needs to be a not-A-world. The closest not-B and not-C-world could be farther away than both the closest not-B-world and the closest not-C-world; hence it may not be a not-A-world. Hence, unlike Dretske's account, the counterfactual account does not satisfy the Conjunction principle. In the counterfactual account a signal that carries the information that B and the information that C separately may not carry the information that B and C. This is counter-intuitive. Since the Conjunction principle is intuitively correct and there is no reason to limit its application, the failure of the counterfactual account to satisfy it marks a disadvantage compared to Dretske's account.

In short, the counterfactual account has the advantage of conforming to the limited application of the Xerox principle, but it leads to an unacceptable consequence with respect to the Conjunction principle. On the other hand, Dretske's theory satisfies the Conjunction principle, but does not match the limited application of the Xerox principle. There is no winner in this game; the result at best is a tie. Hence, it is not true that information flow properties provide reason to prefer the counterfactual account over the Dretskean account as Cohen and Meskin claim.

B. Information and Laws

The second advantage of the counterfactual account, according to Cohen and Meskin, is that it has a less expensive ontology than Dretske's account, which appeals to natural laws. Dretske's theory needs natural laws or nomic dependencies between a signal and its informational content to distinguish genuine information-carrying relations from coincidental correlations. For example, if your room and my room have the same temperature at a given time, the thermometers in both rooms will display the identical reading. Yet it would be wrong to say that the thermometer in your room carries information about my room's temperature. For information-carrying relations, there needs to be a lawful dependency between the reading of the thermometer and the temperature of the room. Whereas this dependency holds between the thermometer in my room and my room's temperature, there is no such relation between my room and the thermometer in your room. Although Dretske's definition does not explicitly mention this nomic dependency, he is clear that assigning unity to the conditional probability results directly from nomic dependencies:

In saying that the conditional probability (given r) of s 's being F is 1, I mean to be saying that there is a nomic (lawful) regularity between these event types, a regularity which *nominally precludes* r 's occurrence when s is not F .

[Dretske 1981: 245 (emphasis in original)]

Thus Dretske recognizes that his conditional probability theory of informational content presupposes the existence of natural laws.

On grounds of ontological economy, an informational account that does not appeal to nomic dependencies is preferable to one that does make such an appeal. Cohen and Meskin claim that their account does not require such lawful regularities. They recognize that counterfactuals are frequently considered to presuppose nomic dependencies between their constituents [cf. Goodman 1954]. But, they continue, this is not ‘untendentious’ and ‘some think it is a mistake to characterize counterfactuals as essentially dependent on laws’ [Cohen and Meskin 2006: 338]. Cohen and Meskin take no stand on this issue, professing to be agnostic about whether counterfactual relations are essentially dependent on the existence of natural laws. Thus they claim that their counterfactual theory is less ontologically committed than Dretske’s probabilistic account.

It is not at all clear, however, that their theory permits them the luxury of their claimed agnosticism. To see why, let us revisit their formulation of the counterfactual definition:

(W) x ’s being F carries information about y ’s being G if the counterfactual conditional ‘if y were not G , then x would not have been F ’ is non-vacuously true.

Although this definition makes no direct reference to nomic dependencies, it is incomplete unless the truth condition of a counterfactual claim is specified. This is because in order to identify an instance of an information-carrying relation one must be able to assess the truth value of the relevant counterfactual claim. Now, specifying truth conditions of counterfactual claims means providing a semantics for counterfactuals. The standard canonical semantics for counterfactuals is the possible worlds semantics. Once such a semantics is introduced, Cohen and Meskin’s claim about ontological economy becomes highly questionable. It is not apparent that an approach requiring possible worlds is ontologically more economical than one requiring natural laws.

In short, although Cohen and Meskin’s counterfactual definition does not refer to a specific ontology, their definition must specify truth conditions for counterfactual claims. Once this is done, the counterfactual view becomes at least as ontologically expensive as Dretske’s probabilistic account. Cohen and Meskin might claim in reply that they want to be agnostic about the truth conditions of counterfactuals. I do not think that such a move is available to them since it would make their account incomplete. Moreover, if they were to take that approach, their theory would be identical to Loewer’s 1983 proposal to define informational content in terms of backtracking conditionals. That is, agnosticism about truth conditions of counterfactuals would make Cohen and Meskin’s account a variation of Loewer’s proposal. In his review of Dretske’s theory Loewer proposes the following definition of informational content:

r ’s being R carries the information that s is F iff r is R and if r is R then s must have been F

[Loewer 1983: 76]

Lewis [1979] calls the conditional claim in this definition ‘a backtracking conditional’. Now Loewer’s proposal makes no reference to laws or nomic dependencies. Natural laws come into play only when Loewer defines truth conditions for backtracking conditionals as follows:

Truth conditions for these, as for other conditionals, are (approximately) ‘there are laws L , conditions C which are co-tenable with $R(r)$ such that $L \& C \& R(r)$ imply $F(s)$ ’.

[Loewer 1983: 76]

If Cohen and Meskin choose not to specify truth conditions for counterfactuals, then the same move must be available to Loewer as well. However, if this is the case there is no significant difference between Cohen and Meskin’s counterfactual account and Loewer’s proposal. We should note that Cohen and Meskin discuss the similarity between their account and Loewer’s backtracking conditional based account. So, pointing out the similarity adds nothing significant to the discussion. However, my point is not only that the theories are similar in this respect, but that maintaining agnosticism with respect to truth conditions of counterfactuals results in an incomplete account of information-carrying relations.

C. Doxastic States and the Naturalism Constraint

In Dretske’s account, the information carried by a signal is relative to the background knowledge of the recipient of the signal. This feature has two motivations. First, it conforms to Shannon’s analysis of information as uncertainty reduction. One’s background knowledge surely determines the amount of uncertainty reduction that a signal provides. If you don’t know that the city of Urfa is located in Turkey, then the signal ‘Hilmi was born in Urfa’ will not reduce your uncertainty about the country where Hilmi was born. However, the same signal for someone who knows that Urfa is in Turkey will completely reduce their uncertainty about the country where Hilmi was born. Second, referring to background knowledge for information-carrying relations matches our common intuitions about information flow.

Cohen and Meskin accept these motivations for referring to background knowledge in the definition of information-carrying relations, but they claim that it is problematic for an objective and naturalistic analysis of mental content. The whole point of using information-theoretic concepts is to provide a naturalistic account of mental content, belief, and knowledge. If the account uses semantic concepts in its definitions, then the result will be a circular and non-naturalistic theory, thereby defeating the purpose. Their worry is that referring to background knowledge makes Dretske’s definition of informational content non-naturalistic and his definition of knowledge circular. But Dretske has a satisfactory answer to both objections. In each case, the semantic reference can be eliminated by backwards iteration. In other words, both of these definitions are recursive definitions. In Dretske’s definition of informational content, the variable k has this recursive

character, and a backwards iteration⁵ will provide the base of the recursion where there is no reference to background knowledge. The same is true of Dretske's definition of knowledge. Dretske says that with a continuing application of his analysis of knowledge and information, 'we reach a point where the information carried does not depend on any prior knowledge about the source, and it is this fact that enables our equation to avoid circularity' [Dretske 1981: 86].

Cohen and Meskin are not convinced by the backwards iteration reply to the circularity objection. They question whether Dretske's recursive definition has a base for all cases, which it needs to succeed. Now the burden of proof clearly lies on Cohen and Meskin to show at least one case where backwards iteration does not stop. They try to provide such an example by exploiting the possibility of two pieces of information mutually depending on each other, as follows:

[L]et it be that, as Dretske claims, K's knowledge that s is F depends on the information that s is F and therefore (because of the role prior knowledge plays in his analysis of information) also on some other bit of knowledge K has about s (e.g., that s is G). For the same reasons, it seems entirely possible that K's knowledge that s is G depends on some further bit of knowledge K has about s. But nothing in Dretske's account rules out the possibility that this further bit of knowledge is in fact K's knowledge that s is F; on the assumption that the dependencies under discussion are transitive, an immediate regress ensues.

[Cohen and Meskin 2006: 341]

At first glance, this example of two mutually dependent pieces of information is problematic for Dretske's theory. However, the purported difficulty becomes murky when one asks under what conditions such a mutual dependence can occur. Suppose the mutual dependence results from either an analytic or a nomic connection between two pieces of information. In both cases, in Dretske's theory, two pieces of information would be carried by the same signal, and so the pieces will be an instance of nested information. In other words, in such cases the mutual dependency that leads to regress does not exist, because these two pieces of information cannot be separated from each other. Dretske develops the notion of nested information precisely to handle such cases:

For if a signal carries the information that s is F, and s's being F carries, in turn, the information that s is G (or t is H), then this same signal also carries the information that s is G (or t is H). For example, if r carries the information

⁵A referee points out that the result of backwards iteration may lead to lack of total objectivity in Dretske's framework, and this problem may well be unavoidable. I think that s/he is right, but as long as the definition of knowledge or informational content can be reduced to basic objective relations, lack of total objectivity is not a problem for Dretske's theory. Moreover, as the referee claims, this could be an unavoidable feature of the phenomena to be explained. S/he provides an analogous situation: Bayesian probability is considered to be a subjectivist account of epistemic probability because it appeals to subjects' prior probabilities before conditionalizing in using Bayes' Theorem. Carnap's confirmation theory, on the other hand, is considered to be objectivist because it starts with the assumption of a probability of 0.5 for all events. However, Bayesians claim that this sort of 'subjectivity' is a virtue over Carnap's confirmation theory because it is an essential feature of the phenomena to be explained.

that *s* is a square, then it also carries the information that *s* is a rectangle. . . . This point may be expressed by saying that if a signal carries the information that *s* is *F*, it also carries all the information *nested in s's* being *F*.

[Dretske 1981: 71 (emphasis in original)]

In short, Cohen and Meskin's example of mutually dependent information does not provide a counterexample to Dretske's claim that the backwards iteration in his recursive definition of informational content can eliminate reference to doxastic states. Before I conclude this section, let me quote what Cohen and Meskin say about possible answers to the circularity objection:

Whether or not Dretske has further apparatus that could block this sort of circularity, the broader point is that Dretske has given us no reasons for believing that his analysis of information ever breaks out of the intentional/doxastic circle.

[Cohen and Meskin 2006: 341]

But Dretske's notion of nested information can block this sort of circularity. If this is not a reason to think that Dretske's analysis of information can break out of the intentional/doxastic circle for the type of cases Cohen and Meskin discuss, it is not clear what could count as such a reason.

IV. Conclusion and Suggestions

Cohen and Meskin claim that their counterfactual account is preferable to Dretske's conditional probability account on three grounds: first, it conforms to the limited application of the Xerox principle; second, it has greater ontological economy; and third, it does not refer to doxastic states. As I have argued, none of these is really an advantage. For the last two claimed advantages, Cohen and Meskin's objections to Dretske's theory are not justified. And although their first objection is justified, their counterfactual account fails to satisfy another essential principle, namely the Conjunction principle.

Given this analysis, there is no good reason for giving up Dretske's probabilistic account in favour of the counterfactual approach. However, we should recall that the main incentive to examine the counterfactual account was another serious problem for Dretske's theory, the lack of a probability interpretation that grounds inverse conditional probabilities. Despite the fact that the alleged advantages of the counterfactual account are not real advantages, it is true that it avoids the probability interpretation problem. Hence, preferring Dretske's definition of information-carrying relations will not be well-justified until we find a solution to the probability interpretation problem. To my knowledge, there has been no attempt to solve this problem within the Dretskean framework. As Cohen and Meskin conclude:

We are left, therefore, without a suitable way of understanding the probabilities that Dretske uses to underpin his theory of information. One response to this situation would be to hope for some new account of probability that avoids these difficulties. For those who, like us, are too impatient to wait for that outcome, the counterfactual account of information will seem attractive, insofar as it sidesteps the problems about probabilities altogether.

[Cohen and Meskin 2006: 337]

Impatience, however, can easily lead to mistakes and patience always pays off. In fact, there are two promising paths for finding a new account of probability that can solve the problem. The first is using an available account; here the relative frequency approach looks to be the best candidate. The second is much more radical: perhaps we have to question a fundamental assumption underlying the interpretation of probability. Let me briefly sketch these two approaches.

We have seen that interpreting probability as the degree of belief leads to a non-naturalistic and subjective account, which is not acceptable for our naturalistic ambitions. Propensity interpretations do not support inverse conditional probabilities. So, there is no point in pursuing these interpretations as solutions to the probability interpretation problem afflicting Dretske's theory. The relative frequency interpretation, however, is more promising. Although it does not apply to singleton events, this weakness is not specific to Dretske's definition. Moreover, the relative frequency approach is commonly used for empirical pursuits. As a result of this prevalence, there have been attempts for solving the problem of singleton events. One of these attempts deserves attention: the Reichenbach-Salmon solution. Reichenbach [1949] suggested using the notion of 'weight' for singleton events instead of the notion of 'probability'. Wesley Salmon [1966] improved Reichenbach's solution with some revisions. Thus, it may be useful to examine the advantages and shortcomings of the Reichenbach-Salmon approach for grounding Dretske's use of inverse conditional probabilities. This is the first promising research path.

The second suggestion is more radical. The main assumption in the debate about interpretations of probability is that there is a tension, if not a diametrical opposition, between subjective and objective interpretations. But perhaps this fundamental assumption is mistaken. One could explore the possibility of denying this assumption and claim that the notion of probability is partly subjective and partly objective. In fact, such a move proved to be fruitful in another context. Daniel Dennett's [1987] answer to the question under what conditions intentionality could be attributed to an organism relies on such a move. He claims that the intentionality we attribute to human beings is a result of a stance that we take, i.e., the intentional stance, which he contrasts to physical and functional stances. Then, in considering the import of the intentional stance, he maintains that it is both objective and subjective. It is objective because it picks out the causal threads in the physical structure of the organism to which intentionality is attributed. It is subjective in so far as it relies on the

explanatory purposes of the entities that attribute intentionality to the organism. Although an analysis of the merits of Dennett's approach is beyond the scope of this paper, my point is that his approach may provide a model for a similar move in the context of interpretations of probability.

In a nutshell, the second research path may go something like the following. Let's assume that an organism's degree of belief of the likelihood of an event's happening is determined by the past experiences of the organism. The degree of belief, in this situation, is subjective because it is determined by the subjective experiences of the organism. On the other hand, the past experiences are a part of the external world, and thus conform to the objective laws of nature. The motivation here is significantly similar to Huw Price's notion of 'subject naturalism' [Price 2004].⁶ Price discusses two different forms of naturalism: subject naturalism versus object naturalism. The context in which he provides the distinction is the relationship between natural science and philosophy. Subject naturalism considers humans as natural creatures, 'and if the claims and ambitions of philosophy conflict with this [scientific view], then philosophy needs to give way'. Moreover, he claims that the perspective of the organism (subject naturalism) is much more fundamental than the no-perspective approach of natural sciences (object naturalism).⁷ Although the context in which this distinction is analysed is different, it is applicable to probability interpretations as well. I think that it is a very promising research project to explore the consequences of applying Price's distinction and Dennett's intentional stance move to the probability literature.

Unfortunately, I do not have enough space to elaborate on these two promising paths for solving the probability interpretation problem. For now, they are just suggestions for future research. The story about them shall be told at another time. They do, however, provide reasons to be more optimistic than are Cohen and Meskin about the success of Dretske's account of informational content.

Bilkent University
California State University, San Bernardino

Received: October 2006
Revised: January 2007

References

- Bremer, Manuel E. 2003. Do Logical Truths Carry Information?, *Minds and Machines* 13: 567–75.
- Carnap, Rudolf and Yehoshua Bar-Hillel 1952. *An Outline of a Theory of Semantic Information*, Technical Report 247, Research Laboratory of Electronics, Cambridge MA: MIT Press.
- Cohen, Jonathan and Aaron Meskin 2006. An Objective Counterfactual Theory of Information, *Australasian Journal of Philosophy* 84: 333–52.
- Demir, Hilmi 2006. *Error Comes with Imagination: A Probabilistic Theory of Mental Content*, Ph.D. Thesis, Bloomington IN: Indiana University.
- Dennett, Daniel 1987. *The Intentional Stance*, Cambridge MA: MIT Press.
- Dretske, Fred 1983. Author's Response, *Behavioral and Brain Sciences* 6: 82–90.
- Dretske, Fred 1988. *Explaining Behavior: Reasons in a World of Causes*, Cambridge MA: MIT Press.
- Dretske, Fred 1981. *Knowledge and the Flow of Information*, Cambridge MA: MIT Press.

⁶I thank the referee who brought Huw Price's subject naturalism to my attention. He pointed out that the literature on pragmatism and probability is also useful in this context.

⁷It is important to note that Price's ideas are relevant to an important distinction in neuroscience literature: the distinction between 'the animal's perspective' and 'the observer's perspective' [Rieke et al. 1999].

- Dretske, Fred 1994. Misrepresentation, in *Mental Representation: A Reader*, ed. S. Stich and T. Warfield, Oxford: Blackwell: 157–70.
- Eliasmith, Chris 2000. *How Neurons Mean: A Neurocomputational Theory of Representational Content*, Ph.D. Thesis, St. Louis MO: Washington University.
- Goodman, Nelson 1954. *Fact, Fiction, and Forecast*, Cambridge MA: Harvard University Press.
- Hintikka, Jaakko 1970. Surface Information and Depth Information, in *Information and Inference*, ed. J. Hintikka and P. Suppes, Dordrecht: Reidel: 263–97.
- Lewis, David 1973. *Counterfactuals*, Oxford: Basil Blackwell.
- Lewis, David 1979. Counterfactual Dependence and Time's Arrow, *Noûs* 13: 455–76.
- Loewer, Barry 1983. Information and Belief, *Behavioral and Brain Sciences* 6: 75–6.
- Loewer, Barry 1987. From Information to Intentionality, *Synthese* 70: 287–317.
- Price, Huw 2004. Naturalism without Representationalism, in *Naturalism in Question*, ed. D. Macarthur and M. De Caro, Cumberland RI: Harvard University Press: 71–88.
- Rieke, F., D. Warland, R. deRuytervanSteveninck, and W. Bialek 1999. *Spikes: Exploring the Neural Code* (Computational Neuroscience), Cambridge MA: MIT Press.
- Reichenbach, Hans 1949. *The Theory of Probability, An Inquiry into the Logical and Mathematical Foundations of the Calculus of Probability*, Berkeley CA: University of California Press.
- Salmon, Wesley 1966. *The Foundations of Scientific Inference*, Pittsburgh: University of Pittsburgh Press.
- Scarantino, Andrea 2005. *Did Dretske Learn the Right Lesson from Shannon's Theory of Information?*, Pittsburgh: University of Pittsburgh Press.
- Shannon, Claude E. 1948. A Mathematical Theory of Communication, *Bell System Technical Journal* 27: 379–423, 623–56.
- Usher, M. 2001. A Statistical Referential Theory of Content: Using Information Theory to Account for Misrepresentation, *Mind and Language* 16: 311–34.