

Trinity University

Digital Commons @ Trinity

Philosophy Faculty Research

Philosophy Department

6-1987

The Causal Indicator Analysis of Knowledge

Steven Luper

Trinity University, sluper@trinity.edu

Follow this and additional works at: https://digitalcommons.trinity.edu/phil_faculty



Part of the [Philosophy Commons](#)

Repository Citation

Luper-Foy, S. (1987). The causal indicator analysis of knowledge. *Philosophy and Phenomenological Research*, 47(4), 563-587. doi:10.2307/2107229

This Article is brought to you for free and open access by the Philosophy Department at Digital Commons @ Trinity. It has been accepted for inclusion in Philosophy Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

The Causal Indicator Analysis of Knowledge

STEVEN LUPER-FOY

Trinity University

In this paper I want to describe and motivate an approach to knowledge that I call the Causal Indicator Analysis. My strategy will be to sketch (in Part I) the main features of an adequate account of knowledge, then use my sketch (in Part II) to reveal some of the faults of some of the main analyses defended today. I will be particularly interested in discussing the work of Fred Dretske, whose views have significantly influenced my own. With these tasks behind me, I will offer my own account in Part III, and argue that it has the features an account of knowledge should have.

I. The Main Contours of Knowledge

About many of the things we believe, we are *correct* but no thanks to the way we arrived at them. It is a sheer *coincidence* that we are correct, just as it is a coincidence that my belief about the next winner of the State Lottery is correct if I base it merely on wishful thinking. No such belief counts as knowledge. Indeed, with but one qualification which I will come to later, knowledge just *is* nonaccidentally correct belief.¹ So if we can get clear about what is and is not involved in nonaccidentally correct belief, we will have the bulk of an analysis of knowledge. I will defend three principles about the nature of a belief that is not accidentally correct. Collectively, these principles can then serve as a characterization of the main contours of an adequate analysis of knowledge. A fourth principle, one that does not lay out a requisite of nonaccidentally correct belief, will then be added to my sketch.

¹ Peter Unger offers as an analysis of *S*'s knowledge that *p* the requirement that *S* be correct in believing that *p* but not accidentally. See his "The Analysis of Factual Knowledge," *The Journal of Philosophy* 65 (1968), pp. 157-70.

1. *First Principle: Minimal Reliability*

Imagine that Flip tosses a coin which lands heads up. He takes this to indicate that a light (out of sight in the next room) is on. Tails he takes to indicate that the light is off. Jessica, meanwhile, covertly switches the light on when the result of Flip's toss is heads, and off when the result is tails. Flip's belief that the light is on is a clear example of a belief that is accidentally correct. The *reason* Flip's belief is accidentally correct, I suggest, is that his belief does not have a sufficiently reliable source.

Contrast an example in which Norm arrives at his belief that there is a table in front of him by using the normal visual sequence in normal conditions. This is a clear example of a belief that is *not* accidentally correct. The reason is that Norm's belief, unlike Flip's, has a sufficiently reliable source.

Because the source of Norm's belief is sufficiently reliable to permit knowledge while the source of Flip's is not, contrasting these two cases will help us learn when it is that a source has the minimal reliability requisite for nonaccidentally correct belief and hence for knowledge. Notice that when used in a wide variety of circumstances, the visual process tends to accurately indicate that beliefs about tables are true when and only when they *are* true. By contrast, coin tossing is very unsuccessful when used as a source for beliefs. Apparently, then, the source of a given belief is (at least) minimally reliable just in case it has a (possibly unexercised) *propensity* to accurately indicate the truth about that sort of belief.

To be reliable, a source would have to perform well in a wide variety of conditions, but it need not perform well in all *possible* conditions. In circumstances in which visual illusions are present, vision performs poorly indeed, but it is still reliable overall. The reason we consider vision reliable in spite of its poor performance in situations where illusions are involved, I suggest, is that circumstances in which illusions are present are relatively abnormal. It is the fact that vision does and would perform well if relied on in *normal* circumstances that leads us to judge it reliable. Normal circumstances relative to an individual are (let us say) those in which the people in his or her community typically find themselves.

Let me sum up these reflections. In order for us to be nonaccidentally correct in thinking that a given belief is true, it must have a source that has at least a minimal level of reliability as an indicator of the truth of that sort of belief. In order for our source to possess the minimal level of reliability, it must have the propensity to accurately indicate the truth of that sort of belief in normal circumstances. Knowing that a given belief is true entails that we are not accidentally correct about it; hence the propensity I just described is requisite to a source capable of generating knowledge. Let us put these conclusions down as Principle (I):

- (I) In order for person *S* to know that a given belief is true, its cause must be reliable to at least a minimal degree, and that cause will be reliable to at least the minimal degree if it has the propensity in normal circumstances to accurately indicate the truth of beliefs of the sort at hand.

2. *Second Principle: Community Relativity*

A second principle will help clarify the first. What constitutes normal conditions, the sort of situation in which members of a community typically are ensconced, will vary from one community to the next. As a consequence, the minimum level of reliability a source must possess fluctuates across a range of communities.

Consider a fanciful example. Imagine that on a distant planet called Twin Earth is a community of people like us who are living in conditions like ours.² There has never been and will never be any substantial contact between us and them. Most of the liquid on their planet consists of a substance (call it 'XYZ') that is qualitatively similar to water, only it has a distinct microcomposition. But there are scattered specimens of bona fide water at various places, some of which the aliens know about.

Such an alien could not arrive at a nonaccidentally correct belief that the stuff he or she just scooped up into a cup is water by the method to which we earthlings typically resort, namely by simply seeing and tasting the stuff. That method is reliable enough for us given the circumstances in which we generally find ourselves, but given the circumstances they are typically in, it is not reliable enough for the aliens. Even if one of the aliens were surrounded for hundreds of miles by water and no XYZ, the method of sight and taste would still not be reliable enough. For nonaccidental correctness is secured by belief sources which are minimally reliable in normal circumstances, and typically the aliens find themselves in the company of a substance that is visually and tactilely indistinguishable from water, namely XYZ.³

² This Twin Earth example is borrowed from Hilary Putnam; see "The Meaning of 'Meaning'," *Mind, Language and Reality (Philosophical Papers, Vol. 2)*, (Cambridge: Cambridge University Press, 1975). pp. 215-72.

³ The situation of this alien vis-à-vis the belief that the stuff just scooped up is water is parallel to Flip's belief that his light is on. The source of Flip's belief is completely unreliable (relative to Flip). Nonetheless, given his circumstances (given the fact that Jessica is set to intervene), it is true to say that if he were to arrive at his belief through the coin toss method, his belief would be true. The reason Flip's belief is nonetheless unreliably arrived at is that circumstances in which coin tossing would generate true beliefs about lights are abnormal for Flip. Like Flip's, the source of the alien's belief is completely unreliable (rel-

A second principle will serve to remind us that what constitutes accidental correctness varies from one community to the next. The conditions for knowledge must reflect this fluctuation.

- (II) The minimal reliability of a knowledge-yielding belief source is community-relative.

3. *Third Principle: Situationally Relevant Possibilities*

No matter *how* reliable a belief source is, it is virtually always possible for it to yield the same type of belief in circumstances such that in them the belief is accidentally correct and hence not known. Given his circumstances, Norm's belief was not accidentally correct. But suppose that Abnor is in ones that are rigged as follows. Though Abnor's belief is caused by seeing a real table, he is surrounded by holograms of tables which are visually indistinguishable from the real thing. Then his belief that there is a table in front of him would be accidentally correct, yet it would have precisely the same source as Norm's had.

The notion of rigged circumstances is important enough to make more precise. The rough notion is that of a situation in which someone's belief is arrived at through a reliable source but the person is at best accidentally correct about the belief. But we can do better. Let us say that a set of circumstances is rigged relative to person *S*, the belief that *p*, and a normally reliable belief source (or sustainer) *P* if and only if in them (a) *P* is the source (or sustainer) of *S*'s belief that *p*, but (b) even though *P* has given rise to (or sustained) *S*'s belief that *p*, *p* is or might be false. Thus (b) asserts that the following subjunctive conditional holds in *S*'s situation:

not-(*S* arrives at (or sustains) the belief that *p* through source *P* \rightarrow *p*).

(Here the arrow is short for the subjunctive 'if-then'.)⁴ Notice that 'normal', on my usage, does not mean 'unrigged'; it is conceivable that rigged circumstances *are* normal for some persons (as in certain brain-in-vat cases).

ative to the alien). However, given the alien's circumstances (given that the alien is surrounded for miles by water and no XYZ), it is quite true to say that if the alien were to arrive at his or her belief by seeing and tasting the stuff in the cup at hand, the belief would be true.

⁴ As I explain later in the paper (Part III), I assume that a conditional of the form '*p* \rightarrow *q*' is true just in case *q* holds throughout those near worlds to the actual world in which *p* holds.

An alternative (but less clear) version of (b) would say that even though *P* has given rise to (or sustained) *S*'s belief that *p*, it is not probable that *p* is true.

It is in the sense just indicated that Abnor's circumstances are rigged. The source of Abnor's belief that a table is before him is (roughly) his coming to be in a state of seeming to see a table before him. But he is up to his neck in holograms. Hence there might well have been no table in front of him even though he believed there was through seeming to see one.

Relative to virtually any (contingent) belief and its source, no matter how reliable that source is, there are rigged circumstances. In such situations, that source is incapable of yielding knowledge. But we must be careful not to draw too strong a conclusion from the fact that virtually any source can fail to yield known beliefs in some situation or another. Abnor and Norm have relied on the same source in arriving at their beliefs, and Abnor has not succeeded in knowing the truth of his belief; but it does not follow that *Norm* has failed. Whether or not people know depends in part on the nature of their actual circumstances. Abnor fails to know because his circumstances are *rigged*. But Norm's are not. In his circumstances, if he were to arrive at his belief the way he did, his belief *would* be true. Hence there is no reason to deny that Norm knows that a table is in front of him even though he has acquired that belief through a source that was incapable of yielding knowledge in Abnor's situation.⁵

In order to know that a given belief is true, then, we must have a source for it that enables us to eliminate some of the possible ways in which our belief would be false, but not all of them. A state of affairs in which our

⁵ It is worth noting that Abnor, like Norm, genuinely *sees* the table in front of him but, unlike Norm, Abnor fails to *know* that there is a table there. The visual process does not necessarily lead to knowledge, as the case of Abnor shows.

Exactly what processes are visual is hard to say; the issue is complicated by the thoughts attributed to David Lewis' pet Bruce LeCatt in "Censored Vision," *Australasian Journal of Philosophy* 60, June 1982, pp. 158-62. According to Lewis' pet, we will arrive at an unacceptable view if we say that people see just in case "the scene before their eyes causes matching visual experience as part of a suitable pattern of counterfactual dependence." The problem is that the desired dependence mediated by a perfectly good visual mechanism between an experience and the scene at hand may be eliminated by a censor who is "ready to see to it that I have precisely that visual experience and no other, whatever the scene may be," so that "so long as the scene is such as to cause the right experience, the censor does nothing," while "if the scene were any different the censor would intervene and cause the same experience by other means."

Because seeing and knowing are not unrelated, and because I eventually offer a subjunctive analysis of knowing, the issue arises as to whether my analysis is vulnerable to counterexamples involving such Catesian censors. But my dog Moose assures me that this concern is unfounded. While it is true that we may see things even in the presence of Catesian censors, we do *not* know that the corresponding beliefs are true. For whether I get these beliefs right depends on whether my scenery happens to correspond to a series of experiences of a sort that a nearby censor wishes me to have; if my scenery goes differently in *any* way, my visual beliefs will be mistaken.

belief would be false is an *alternative* to our belief; a knowledge-yielding belief source is one such that, given our circumstances, no alternatives to our belief would hold were we to arrive at our belief through that source. An alternative that, given our situation, *might* arise were we to rely on our source is a *relevant* alternative relative to our situation, our belief, and its source; my suggestion, then, is that we are situated as we must be if we are to know that a given belief is true only when there is no relevant alternative relative to our situation, our belief, and its source.

There is an important sense in which one set of circumstances can be *more* rigged than another. Suppose that when we compare one set of circumstances to another we find that in the first there are more ways in which our belief might turn out to be false if we were to arrive at it through a given source. (That is, there are more relevant alternatives relative to the first situation, our belief, and its source.) Then the first set of circumstances is more rigged than the other (relative to that belief and its source).

When we are situated in certain circumstances, it is possible to know the truth of a given belief only if we rely on a source which is even *more* reliable than one that permits us to know the truth of the same belief in another situation. For in some situations, great reliability is required in order for a potential source to prevent relevant alternatives from arising. Greater reliability is required in a new situation when (though not only when) all of the alternatives which had to be ruled out in the old situation must still be ruled out, and new alternatives must be ruled out as well.

Consider an example. Gemologists presumably have foolproof tests for identifying stones as diamonds in the everyday circumstances with which they are confronted; in such circumstances, no alternatives to their belief that a given stone is a diamond arise which cannot be ruled out by the best techniques. But suppose that in an obscure part of Africa are stones which any gemologist would take to be diamonds even though these gems have an atomic structure that is significantly unlike that of diamonds. A gemologist who comes upon the only real diamond there is in no position to know of the lone genuine stone that it is a diamond. For the possibility that a look-alike is at hand is one that might arise if the gemologist were to use ordinary techniques; so only if the expert used methods capable of distinguishing diamonds from the look-alikes would it be possible to arrive at knowledge in these circumstances. Supposing also that all of the alternatives that arise and must be dealt with in everyday circumstances arise and must be dealt with in this region of Africa as well, then knowledge-yielding techniques in the presence of the look-alikes must be more reliable than usual.

Again, however, we must avoid drawing excessively strong conclusions from facts about the requirements for knowing in special circumstances like the gemologist's. It is true that an extremely reliable diamond detection test must be used in order to arrive at knowledge when we are surrounded by very convincing look-alikes. But extra reliability is unnecessary when we are *not* in such a situation. Coincidentally correct belief can be avoided without it. How reliable knowledge-yielding sources must be *varies* depending on the nature of our circumstances. The minimal reliability is set in normal conditions: a source that is reliable enough to eliminate accident there yields knowledge there. But in other circumstances the minimum might not be sufficient. As the case of the gemologist shows, a *sliding scale* of reliability is required for knowledge.

The African gemologist's circumstances were rigged to a greater degree than are the circumstances which most gemologists face, and we just saw that the African gemologist's sources have to be especially reliable in order to yield knowledge about nearby diamonds. It is tempting to conclude that knowledge-yielding belief sources must always be especially reliable when relied on in situations that are rigged, but that would be a mistake. Consider the situation in which Abnor finds himself. His circumstances are rigged relative to his belief and its source, while Norm's are not; in order for Abnor to arrive at knowledge in his circumstance, he must switch his allegiance to a different source, one such that the possibility that he is confronting a hologram is not a relevant alternative. But he could do so without relying on vision at all and instead by reaching out and carefully feeling the object in front of him to make sure it has the shape of a table. And as far as I can tell, this tactile process is not more reliable than the visual process Norm used. People can be fooled by the felt shape of a table when the object before them is a papier-mâché replica; tactile illusions may be as common as visual illusions. Whether or not the two processes differ in the degree to which they are reliable sources for beliefs about tables is irrelevant, however, so long as both possess the minimal reliability discussed earlier. Beyond that, a knowledge-yielding source must enable us to eliminate the possible ways in which, given the situation, our belief might be false. And in many situations a method which can do that need not be more reliable than one which cannot.

Let us now gather the observations we have made into a third principle, one which spells out what is required of a knowledge-yielding source over and above minimal reliability:

- (III) No belief is known to be true by us if we are in circumstances that are rigged relative to that belief and its source.

4. *Fourth Principle: Against Irrationality*

A fourth principle should be mentioned since it will play a crucial role in the proceedings below. Unlike the previous three, however, this requirement is not motivated by considerations about the requisites of nonaccidental correctness — or so I will indicate in Section III.

It is obvious that irrationality and knowledge are irreconcilably reluctant bed fellows. Any account of knowledge must therefore conform to a final condition:

- (IV) No irrational belief is known to be true.⁶

II. Previous Analyses

Each of these principles has been violated by respectable analyses of knowledge. It will be instructive to survey some of the main theories of knowledge and show why they conflict with our sketch. We will then be in a better position to avoid repeating history.

1. *Violations of the Third Principle*

The third condition, which bans knowledge in rigged circumstances, has been transgressed repeatedly. On one natural understanding, The venerable Traditional Theory,⁷ which equates knowledge with justified, true belief, is a good example. According to an 'internalist' conception of justification, so long as my beliefs and sensory information remain constant as I shift from one situation to another, so will the evidential status of my beliefs.⁸ Let us say that two situations are *doxastically* and *perceptually similar* for a person *S* so long as *S*'s beliefs and sensory information are the same in both. Then on the 'internalist' view, one of my beliefs is justified in a situation only if justified in all doxastically and perceptually similar situations. But on this view, anytime I shift from a situation in which I have a justified, true belief to a doxastically and perceptually similar situation in which my belief is true, my belief remains justified. It does not matter how rigged the new situation is. But clearly the truth of my

⁶ This fourfold characterization of an adequate analysis of knowledge is hinted at in my paper, "The Epistemic Predicament," *Australasian Journal of Philosophy* 62, No. 1, March 1984, pp. 31-34.

⁷ Attributable to A. J. Ayer (see *The Problem of Knowledge* [London: MacMillan Press, 1956], p. 34), among others.

⁸ I argue for a (relatively) 'internalist' conception of justification in my papers "The Reliabilist Theory of Rational Belief," *The Monist*, April 1985, vol. 68, No. 2, and "Surviving Hume's Day," unpublished.

belief need not remain *known* after I shift to rigged circumstances. Standard methods enable gemologist who are situated in normal, unrigged conditions to know diamonds when they encounter them. In such circumstances, these experts will also be justified in believing that they are confronted by diamonds. Nonetheless, their beliefs would remain justified but not known if, without changing their beliefs or sensory information, they were to shift from normal circumstances to ones that are replete with diamond look-alikes. The Traditional Theory, internalistically conceived, cannot succeed because a situation that is *not* rigged relative to a true belief and its evidential basis may be doxastically and perceptually similar to a situation that *is* rigged relative to exactly the same true belief and source.

Causal theories⁹ of various sorts fail for analogous reasons. Such accounts require that my belief that *p* be caused by the fact that *p*, or that there be a law-like connection between these.¹⁰ But my source may continue to meet such causal conditions no matter how rigged my circumstances come to be. For example, Norm's and Abnor's beliefs, which are produced by the same source, both meet the causal condition even though the one belief is generated in rigged circumstances while the other is not.¹¹

⁹ Notable examples include John Locke, *An Essay Concerning Human Understanding*, ed. P. Nidditch (Oxford: Clarendon Press, 1975), Book IV, Chapter XI, Sections 1-2; and Alvin Goldman, "A Causal Theory of Knowing," *Journal of Philosophy* 64, 12 (1967), pp. 355-72.

¹⁰ This type of causal theory is offered by D. M. Armstrong, *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press, 1973), especially pp. 166-71. And as I indicate below, a preliminary version of the account developed by Fred Dretske in *Knowledge and the Flow of Information*, op. cit., is such a causal account.

¹¹ It is interesting to note that our third principle is inconsistent with any account of knowledge which demands of knowledge-yielding belief sources a *static* degree of reliability. (Call people who limit themselves to static conditions 'staticists'.) As we saw earlier when we considered the case of the gemologist, in some situations especially great reliability is required of knowledge-yielding sources, but not in other situations. For this reason, staticists are forced toward one or the other of two unacceptable extremes. They must require that our sources be reliable enough to eliminate accident, but when we are in circumstances like those of the gemologist in Africa, our sources must be terribly reliable to eliminate accident. Thus if we limit ourselves to an analysis that requires just one degree of reliability, we must adopt one that requires a very *high* degree. Unfortunately, this would then make it impossible to know beliefs whose sources are merely reliable enough to rule out accident in unrigged conditions. Consequently we would know very little — which is what the skeptic has always insisted.

There is another option for staticists. They can reduce the level of reliability required of belief sources so that knowledge can be acquired through ones that are only reliable enough to eliminate accident in unrigged circumstances. But the unacceptable consequence of doing so is obvious. There will be beliefs whose sources possess the requisite

2. Violations of the First and Fourth Principles

Let us now consider an analysis which violates our first and fourth principles, the demand that the reliability of belief sources not plunge below a certain minimum, and the demand that our belief not be irrational, respectively. The account I wish to discuss was proposed by Fred Dretske.¹² An analysis offered by Robert Nozick¹³ bears a strong resemblance to Dretske's and suffers from similar flaws. Dretske's simpler analysis can be used to illustrate some points about both views.

The root idea developed by Dretske is this. Suppose that the immediate cause of Norm's belief that there is a table in front of him is his present state of seeming to see such a table. According to Dretske, Norm would then know that a table is in front of him since the following subjunctive conditional is true:

If it were *false* that there is such a table, then Norm would not seem to see one.

More generally: person *S* knows that *p* just in case *S*'s belief has an immediate cause *C* which is a *conclusive reason* for *p*, where *C* is a conclusive reason for *p* if and only if:

$\text{not-}p \rightarrow \text{not-}(C \text{ occurs or holds}).$ ¹⁴

minimal reliability (like those of the gemologist in Africa) and hence count as known even though those beliefs are accidentally correct.

Clearly the only way out of the staticist's dilemma is to shift to *variable* restrictions: ones that require a scale of reliability that often slides upward as the degree to which circumstances are rigged increases.

¹² In "Conclusive Reasons," *Australasian Journal of Philosophy* 49, 1971, no. 1, pp. 1-22.

¹³ *Philosophical Explanations* (Cambridge: Harvard Press, 1981), Section III.

¹⁴ This restriction is very similar to the third condition of Nozick's "tracking" analysis.

I should mention that I settled on my reading of Dretske only after sorting through several obscurities in his account. Let me discuss these briefly.

According to him, "R is a conclusive reason for P" if and only if the following condition holds:

(a) R would not be the case unless P were the case.

But there are two ways to interpret (a):

(b) If P were not the case, R would not be the case.

(c) If R were the case, then P would be the case.

Note that (b) is the contrapositive of (c), but the contrapositives of counterfactual conditionals are not equivalent. I have assumed that (b) rather than (c), (and rather than (b)

Since he is in normal circumstances, Norm's seeming to see a table in front of him is a conclusive reason for his belief that one is there. Given his circumstances, in no worlds close to the actual world does he seem to see a table though none is in front of him. Yet if he were in rigged conditions like Abnor's, he would not have a conclusive reason for his belief. Abnor is surrounded by table holograms; hence there are worlds very close to the actual world in which he seems to see a table though none is there.

Although Dretske's account handles the cases of Abnor and Norm, it performs less impressively when applied to other cases. Earlier I mentioned an example in which Flip comes to believe that a light in the room near him is on by tossing a coin. His belief is clearly irrational, and its source is clearly not reliable enough. Yet on Dretske's analysis, Flip has arrived at knowledge.

Let us say that the immediate cause of Flip's belief is his seeming to see his coin land heads-up. Normally this state of affairs would not be a conclusive reason for a belief about a light in a room, but Flip is not in normal circumstances. Because Jessica is making sure that the light is on whenever

and (c) express Dretske's intention.

There is another interpretive difficulty that must be discussed. At one point in "Conclusive Reasons" Dretske says that (a) means

(e) Whenever (a state such as) R in circumstances C then (a state such as) P,

where C is "those circumstances which are logically and causally independent of the state of affairs expressed by P." And on this understanding of subjunctive conditionals, the Jessica case is not a counterexample to his analysis. In order for Flip to know that the light is on, the following conditional would have to hold:

(f) Whenever a coin toss lands heads up in Flip's circumstances C then the light is on.

But since facts about Jessica are not causally independent of the light being on, they cannot be included in C; (f) is, as a result, false.

It would be possible to revise the Jessica case so as to question Dretske's analysis even on his understanding of counter-factuals. However, I will limit myself to pointing out that Dretske's understanding of subjunctives is wildly implausible. On it, it is false to say that if I were to hit my brakes, my car would stop — even if my brakes are fully operative and conditions are optimal. For the *fact* that my brakes are operative is not causally independent of the fact that my car will stop and hence cannot be part of C in the conditional,

(g) Whenever I hit my brakes in C then my car stops.

The importance of this problem is increased by the fact that in normal circumstances we really *ought to be able* to know that our car will stop on the basis of the fact that we hit our brakes.

Flip's coin lands heads-up, there are no worlds close to the actual world in which Flip seems to see his coin land heads-up yet the lights fail to be on. (If you think there *are* such worlds since human beings such as Jessica are fallible, then replace her with an extremely reliable machine.) Flip possesses the conclusive reason Dretske requires for knowledge.

This coin toss example illustrates an important point. So far we have used the term 'rigged' to refer to circumstances in which we would not, using our admittedly reliable methods, get our belief right; thus we have had in mind a sense in which one situation may be *unfavorably* rigged. As the coin toss example shows, however, our circumstances can also be *favorably* rigged in the sense that in them we would not get our belief *wrong* even though we are using an *unreliable* method.¹⁵ Flip's falls into the second category. Favorably rigged circumstances present a problem because for each belief source, no matter how unreliable, there will be circumstances that compensate for its great unreliability, making it as unlikely as you like that we will get things wrong. But an unreliable source does not yield knowledge even if propped up in this way. Benevolent Cartesian demons cause almost as much trouble as malevolent ones. We can therefore reject any analysis according to which the great likelihood (or certitude) that our belief is correct given its source and given our circumstances is sufficient for knowledge.

I have been discussing an early view of Dretske's; let me pause to consider his latest effort. In *Knowledge and the Flow of Information*¹⁶ he aims to provide an information-theoretic analysis of knowledge. I know that *p*, Dretske says, just in case my belief is caused (or sustained) by an event or state of affairs that carries the information that *p*.¹⁷ He then offers two accounts of when it is that information is carried by a state of affairs, the first of which is supposed to be an approximation later improved on by the second.

The first of Dretske's descriptions of informational content yields a causal analysis of knowledge. State of affairs *r* carries the information that *p*, on Dretske's first view, just in case there is a "lawful (exceptionless) dependence" between *r*-like situations (or events) and *p*-type states of

¹⁵ Our earlier definition of rigged circumstances is actually suitable only as an account of unfavorably rigged conditions. But we can define favorably rigged conditions analogously. A set of circumstances is *favorably* rigged relative to person *S*, belief *p*, and a normally *unreliable* belief source (or sustainer) *P* if and only if in them (a) *P* is the source (or sustainer) of *S*'s belief that *p*, and (b) *p would* be true if *S* were to arrive at (or sustain) the belief that *p* through *P*.

¹⁶ (Cambridge: MIT Press, 1981).

¹⁷ *Ibid.*, p. 86.

affairs.¹⁸ Given the way Dretske explicates knowledge in terms of information, this picture of information commits him to the view that a person *S* knows that *p* just in case there is a nomic regularity between states of affairs of the sort that caused *S*'s belief, and *p*-type states of affairs. We have already seen that causal analyses are not strong enough to meet the third principle of our preliminary sketch of knowledge; Dretske's version is no exception. It falls prey to the same sort of problem which crippled his early conclusive reasons account of knowledge.

Dretske offers the causal version of his analysis in Chapter 4; he refines it in the next. "To know, or to have received information," he says in Chapter 5, "is to have eliminated *all relevant alternative possibilities*."¹⁹ On the new account, state of affairs *r* carries the information that *p* just in case the fact that *r* holds eliminates all possibilities that are both "relevant" and inconsistent with the truth of *p*. And I know that *p* just in case my belief has a cause which carries the information that *p*, i.e., one that eliminates all alternatives to *p* that are "relevant." On the old analysis, there had to be an exceptionless dependence between *r*-type situations and *p*-type situations in order for *r* to carry the information that *p*; the new account is less demanding. Since *r*'s holding need only eliminate the alternatives to *p* that are "relevant," then the dependence between *r*- and *p*-type situations need not hold in the presence of "*irrelevant*" alternatives to *p*.

Dretske's new view seems to be a weakened version of the earlier causal approximation, one designed to allow exceptions to exist in the dependence between *r*-like situations and *p*-type situations whenever the exceptions are "irrelevant." If so, if his new view in no way strengthens his old one, it cannot succeed, for the earlier view is already too weak. But it is hard to judge. The problem is that Dretske's account is useless until we are told when an alternative is relevant. Dretske does make helpful comments about the conditions under which an alternative is relevant. Thus he points out that not all possibilities that are "consistently imaginable" are relevant ones.²⁰ And he claims that what counts as relevant will depend on "contextual," "social" and "pragmatic" factors.²¹ But Dretske's comments are not specific enough to amount to a genuine account of relevant alternatives.

¹⁸ *Ibid.*, note 1 from p. 65, on p. 245.

¹⁹ *Ibid.*, p. 133.

²⁰ *Ibid.*, p. 130.

²¹ *Ibid.*, pp. 132-33.

Moreover, remarks Dretske makes about “channels” suggest that his analysis of information is even too weak to meet the first principle of our sketch of knowledge. He defines a channel of communication as “the set of existing conditions that have no relevant alternative states”²² and on which the purported information carrier (which he calls a ‘signal’) depends. These channel conditions, he says, are to be held fixed when we determine whether the dependency relations crucial to information carriage hold. Unfortunately, it is not clear when a condition has a “relevant alternative state;” but if the example Dretske gives in a footnote is any clue, then holding these conditions fixed makes it possible for us to arrive at knowledge through completely unreliable means.²³ There he describes a factory worker who can pick out resistors from everything else in the factory (which, depending on the factory, could be extremely easy), but who would readily confuse them with things he might see outside the factory. Dretske says that “if . . . we include *his being in the factory* as one of the channel conditions . . . , there is nothing to prevent us from saying that (in the factory) he is getting the information that [a given object] is a resistor.” Yet the worker’s method of picking out capacitors might be completely unreliable, as emerges when we consider its performance in the world outside of the factory.

As far as I can tell, then, Dretske’s account runs afoul of either our first or third principle of knowledge.

3. *Violations of the Second Principle*

Defeasibilist analyses of knowledge have trouble meeting our second condition, which introduced community relativity.²⁴ This difficulty, though not insurmountable, is worth discussing briefly.

The different versions of Defeasibilism are legion, but the intuitive idea is simple. I know that p just in case my belief is both true and adequately justified, and my justification is indefeasible; i.e., to my evidence can be added any true statement and the result would still adequately justify my belief. While Norm meets these conditions, Abnor does not. The statement that there are holograms of tables in his room will *defeat* the visual evidence Abnor has for his belief. Were this statement added to his evidence, he would no longer be adequately justified in his belief about the table. Similar reasoning allows Defeasibilists to say that gemologists may know about diamonds using the usual tests even though these tests would

²² Ibid., p. 123 (see also p. 115).

²³ Note 17 from p. 131.

²⁴ Several versions of the Defeasibility analysis are gathered in *Essays on Knowledge and Justification* (Ithaca: Cornell University Press, 1978).

fail in special circumstances. So far, so good.

Nonetheless there is a difficulty confronting Defeasibility analyses: the problem of misleading defeaters.²⁵ The example discussed earlier of the aliens on Twin Earth neatly illustrates the point that some truths defeat evidence for beliefs we nonetheless *do* know. Suppose that I have arrived at the belief that the stuff in my cup is water by seeing and tasting it. I have good reason to think my belief is true and it is. But suppose that the Twin Earth tale were actually true. Then thousands of people have experienced the very same sort of tastes and visual information I just had yet what they tasted and saw was not water. Clearly this statement defeats my evidence for my belief about my cup of water. Yet the existence of Twin Earth in no way prevents me (or other members of my community) from knowing that this belief is true. The Defeasibility approach is clearly in danger of violating our second adequacy condition.

Those true statements which defeat grounds that are clearly a sufficient basis for knowledge are in some sense *misleading*. If we could specify when it is that a defeater is misleading, then we could repair Defeasibilism by requiring that our justification for a belief need only withstand any putative defeaters that are *not* misleading. An adequate account of misleading defeaters would allow us to say why the evidence Abnor has for his belief concerning his table is genuinely defeated by the statement that his room is filled with table holograms, while the evidence I have for my belief that I am drinking water is only misleadingly defeated by the statement that thousands of people with the same tactile and visual information as I have are not confronted with water. I have not encountered a successful account of misleading defeaters; nonetheless, I lack any proof that none will be forthcoming. Attention to our second principle of knowledge may help Defeasibilists develop such an account.

This concludes our discussion of how several popular approaches to knowledge fail to conform to the demands we sketched in Part I. Let me now provide an account of knowledge which fares somewhat better.

III. The Causal Indicator Analysis

I know that a belief is true just in case my belief is not irrational and my being correct about it is not simply fortuitous. At least, this much is true of beliefs about contingent truths; the analysis of our knowledge of necessary truths I leave aside. But there are two ways a belief might avoid being

²⁵ This problem has been discussed by most Defeasibilists. See e.g., Lehrer and Paxson, "Knowledge: Undefeated Justified True Belief," in *Essays on Knowledge and Justification*, op. cit. Their analysis, like many other versions of Defeasibilism, does not handle the Twin Earth case.

irrational. One is for it to be *nonrational*. This is the attribute any belief has that is not inferentially based on another belief. The second way a belief could fail to be irrational is more familiar: it might be rationally based on other beliefs. All inferential beliefs are either rational or irrational; they are never nonrational.

Because of the relative ease with which noninferential beliefs avoid being irrational, the condition under which they are known to be true is relatively simple. It is sufficient (and necessary) that we be non-accidentally correct about such beliefs. Call this condition the *metaphysical* component of knowledge. But inferential knowledge is slightly more complex. It has a second component: the conditions under which a belief is rational. Call this the *doxastic component*. I will now describe these components in more detail.²⁶

1. *Noninferential Knowledge*

I know of a noninferential belief that it is true just in case I am nonaccidentally correct about it. For this to be possible, my belief must have a source (or sustainer) which has minimal reliability, and relative to my belief and its source my circumstances must not be rigged. Let us consider these in turn.

A belief's source has the minimal reliability requisite for knowledge just in case it has the following feature: usually if that type of belief were acquired through that type of source, then it would be true. But this idea will take some spelling out.

Let me begin with a simplifying assumption. I will presume that matters involving causation can be dealt with in terms of facts, so that there is no need to introduce the ontological category of events. Thus I will avoid saying that beliefs have as their sources events (or event sequences) that caused those beliefs. Rather, I will say that facts cause, and are sources of, beliefs. Typically entire causal *chains* are causes of, and hence sources of, our beliefs. But I will not say that the links of these chains are events which cause their successors and in turn are caused by their predecessors. Instead, these links are facts.

I will need to speak of the *type* of a given fact *f*. The relevant type is given by a certain sort of state of affairs that is involved in that fact *f*. The

²⁶ Let me offer a concession to those who would rather say that the truth of nonrational beliefs cannot be known. The case for the existence of nonrational beliefs that are known to be true consists largely in the intuition that some noninferential beliefs can be known. Those who do not share this intuition can nonetheless adopt the most important aspects of the theory I will offer. Simply adopt my analysis of inferential knowledge as the whole story about knowledge.

relevant state of affairs is arrived at by abstracting from the time at which the fact holds. It will be useful to introduce a device that marks off this salient state of affairs. Thus let f^* be the state of affairs that remains when all of the particular times involved in fact f are omitted. If no such times exist, then f^* just is f .

Some examples will help. Suppose that f_1 is the fact that the cat ate fish at noon, New Year's Day, 1984. Then f_1^* is the state of affairs of the cat's eating fish. Where f_2 is the fact that officer O'Malley ticketed Frieda at noon, New Year's Day, 1984, while the cat ate fish, f_2^* refers to the situation in which officer O'Malley tickets Frieda while the cat eats fish.

I can now state a preliminary version of the analysis of minimal reliability. In order for person S to know that p , one of the causes of S 's belief that p must be a fact (call it f) such that, on most occasions t , if the temporally abstracted version of f held at t , then the temporally abstracted version of p would too:

There is a cause (or sustainer) f of S 's belief that p , and at most times t , the following conditions holds:

$$f^* \text{ at } t \rightarrow p^* \text{ at } t.$$

When

$$f^* \text{ at } t \rightarrow p^* \text{ at } t$$

holds at most times t , let us say that fact f is a *reliable indicator* that p .

Progress has been made; we can now explain why Flip cannot, on the basis of a coin toss, come to know that the lights in the room near him are on. Flip's belief is caused by the fact that f_1 : his coin landed heads-up at a particular time t_1 . f_1^* will be the situation in which Flip's coin lands heads up. Now, typically were Flip (or anyone else) to toss a coin in such a way that it lands heads up, lights in a nearby room might be off. Hence it is false that usually the lights would be on if Flip's coin were to land heads up. The source of his belief does not possess the requisite reliability for knowledge.

For simplicity, I have formulated the analysis of minimal reliability in a way that presupposes that if f is a reliable basis for the belief that p , then p^* would usually be true (at a given time) if f^* were to hold (at that time). The fact that Riley is eating dinner at a given time, for instance, is a good basis for the belief that Riley is alive at that time. But of course this presupposition is not generally true. The following fact

f_1 : Robust Riley, who lived well over 100 years, was ten on New Year's Day, 1810

is good grounds for believing that q : Riley was 100 years old 90 years after New Year's, 1810. However, there *are* no times t such that if Riley were ten at t , then she would be 100 at t . The problem (of nonsynchronicity) is easily enough dealt with, but the solution requires introducing considerable complication into our analysis. I will sketch the solution, then, for the sake of simplicity, omit the new details in later discussion.

Fact q involves a time which is specified in terms of the time at which f_1 says that f_1^* holds. To say that q is to say that q^* holds 90 years after the time f_1 says that f_1^* holds; it is to say that q^* holds 90 years after New Year's, 1810. (Here q^* is the state of affairs of Riley's being 100 years old.) Hence q involves a temporal relation which specifies a time in terms of the time when state of affairs f^* holds. I will need to talk about such temporal relations which specify times in terms of when some states of affairs hold. To do so, let me introduce a definition. Let

$$\text{TR}(t', t'', f, p)$$

refer to that temporal relation involved in the fact that p which specifies how the time t'' at which p^* holds is related to the time t' at which f^* holds. Thus, for example, $\text{TR}(\text{New Year's } 1810, \text{New Year's } 1900, f_1, q)$ is the relation, t' is 90 years before t'' .

The following analysis solves the problem of nonsynchronicity:

There is a cause (or sustainer) f of S 's belief that p , and at most times t , the following condition holds:

$$f^* \text{ at } t \rightarrow p^* \text{ at time } t' \text{ such that} \\ \text{TR}(t, t', f, p).$$

Nothing further will be said about the problem of nonsynchronicity; but other necessary modifications are pressing.

Typically, an entire causal *chain*, each link of which causes the next, will cause a belief. Now, it is entirely possible for some links in a causal chain to be reliable indicators while others are not. So long as one link is a reliable indicator, however, the entire chain as a whole will be as well, according to our account. That is, the fact that the entire chain held will be a reliable indicator since it includes a component which is a reliable indicator. Unfortunately, many beliefs arrived at through clearly unreliable means are produced by chains whose status as reliable indicators is due to their having a reliable indicator as a component.

Here is an example. The fact that a piece of cheese just fell onto the floor of Fran's apartment startles a mouse, which in turn causes it to scurry out the back door. Fran subsequently arrives at the belief that there is cheese in her apartment at noon that day by inferring it from her belief that a mouse scurried out the back door. The mouse's having exited is not a reli-

able indicator that there is cheese in Fran's apartment during that particular noon (or any other). But the fact that cheese fell onto her floor at noon is, and any chain which included it will be too. One such chain caused Fran's belief about the presence of cheese in her apartment. Consequently, her belief has a reliable source according to our analysis.

Two factors in combination explain our hesitation to say that Fran's belief has a genuinely reliable source. First, it is produced by a chain with several links which are unreliable indicators. Second, Fran's *access* to the one link which *is* a reliable indicator is through those links which are not. The reliability of the chain which produced her belief is due to its including the fact that cheese fell onto the floor during one fateful noon-time. But this fact is able to affect Fran's belief only by causing other facts such as the fact that the mouse was startled, and *these* facts are *not* reliable indicators that her belief is true.

The upshot is that we must strengthen our account. To guarantee that a person *S*'s belief that *p* has a reliable source, more must be required than that one of its causes be a reliable indicator that *p*. A causal *chain* can be a reliable source for a belief only if *every link* in it is a reliable indicator. We must therefore supplement our analysis with the requirement that there be a causal chain (possibly with only one link) which produced (or sustains) *S*'s belief that *p* each of whose links is a reliable indicator that *p*.

To say that a chain with reliable links must be behind my belief is not to say that every *part* of those links must be reliable indicators. It will rarely happen that the parts of the links are reliable indicators. One link in a causal chain that produced my belief that there is a pad of paper in front of me is the fact that my retinas have just been stimulated in a given way while I am in such and such an intellectual state. In some sense this link has as a part the fact that a given cell in my left retina has just entered a certain sort of state. While the link itself is a reliable indicator, this part of it is not, since it is highly likely that it enters that sort of state when I am confronted with objects other than pads that bear some resemblance to pads.

In setting out the contours of an adequate account in the previous section I warned against imposing restrictions which are not community bound and hence in violation of Principle (II). Unfortunately, our account of minimal reliability does just that. For imagine that on several distant planets there are people like us living in conditions like ours, and that these planets are densely populated, so that far more of the actual people in existence live on them rather than on Earth. However, unnoticed by the aliens, the water (or most of it) on their planets is somehow transformed into, or replaced with, XYZ. Because of the transformation, the aliens are after a while exposed to XYZ more often than to H₂O. According to our

account of reliability, the existence of such aliens entails that even Earthlings would not normally have a reliable source for their beliefs about water. For among the links of the causal chain through which I typically arrive at my belief that the stuff I am drinking is water, is the fact that I have sensory stimulations of a sort normally associated with water. Clearly, the aliens would usually be confronted with XYZ instead of water if they had such stimulations. Because most of the occasions on which people have those stimulations are occasions on which the aliens have them then, by our analysis, the fact that I have them is not a reliable indicator that the stuff I am drinking is water.

Whether or not I have a reliable source for my belief that that stuff I am drinking is water should not depend on whether such aliens exist or on the fact that stuff on their planet is indistinguishable (using our methods) from water. Our methods need only enable us to handle the sorts of circumstances in which those in our community normally find themselves.

These considerations show that we must specify minimal reliability relative to communities. The following analysis does so:

There is a causal chain which produces (or sustains) *S*'s belief that *p*, and each link of it is a fact *f* such that at most times *t*, the following condition holds:

A member of *S*'s community is caused to believe that p^* at *t* by the fact that f^* holds at $t \rightarrow p^*$ at *t*.

The aliens and Earthlings normally have the same source for their beliefs about water. But according to the new account, that source is reliable for us even though unreliable for them.

We must pause to alter the definition of 'reliable indicator' in light of our new account. We must now consider fact *f* to be a reliable indicator that *p* for *S* just in case the following condition holds at most times *t*:

a member of *S*'s community is caused to believe that p^* at *t* by the fact that f^* holds at $t \rightarrow p^*$ at *t*.

If my belief that *p* has as its source a causal chain each link of which is a reliable indicator that *p*, then that source has the minimal level of reliability necessary for me to know that *p*. But more is necessary for (noninferential) knowledge than minimal reliability. In order to meet Principle (III), our analysis of knowledge must also entail that beliefs cannot be known to be true in circumstances that are rigged.

Fortunately, we can design a condition that accommodates Principle (III) by simply applying our definition of rigged circumstances. Earlier we said that the circumstances in which a belief is produced through a source

are rigged just in case in them that belief might be false even though it is produced through that source. The condition we need simply demands that the circumstances in which we have arrived at a belief *not* be rigged relative to it and its source. More formally:

The source of *S*'s belief is a causal chain each link of which is a fact *f* such that:

$$f \text{ holds} \rightarrow p.$$

Let us say that *f* is a *conclusive indicator* that *p* if and only if:

$$f \text{ holds} \rightarrow p.$$

Following a suggestion made by Nozick, I will understand the subjunctive conditional (here abbreviated with an arrow) in a way that is slightly stronger than usual. It will not do to say that '*f* holds \rightarrow *p*' is true so long as *p* and '*f* holds' are true in the actual world; that makes subjunctives too weak. Instead, I will say that '*f* holds \rightarrow *p*' is true just in case *p* holds throughout those near worlds to the actual world in which *f* holds; that is, *p* must hold throughout the (*f* holds)-*neighborhood* of the actual world.²⁷

It is now possible to explain why Abnor fails to know there is a table in front of him in spite of the fact that the source of his belief is minimally reliable. One link in the causal chain(s) that produced Abnor's belief is his being in a state of seeming to see a table in front of him. Given his circumstances, clearly his belief might have been false even though he is in this state. For there are very close worlds to the actual world in which he is in that state as a result of his having spotted the table hologram.

Let me now summarize the account of noninferential knowledge at which we have arrive. A person *S* knows the truth of *S*'s noninferential belief that *p* just in case:

²⁷ Nozick's suggestion appears in footnote 8 on p. 174, *Philosophical Explanations*, op. cit.

The truth conditions for subjunctives could be clearer, of course. My own inclination is to think that something like the Lewis-Stalnaker approach is correct. Even if this approach fails, however, my account should be of interest since it is compatible with a wide range of semantics for subjunctive conditionals. For the Lewis-Stalnaker view, see David Lewis, *Counterfactuals* (Princeton, 1973), and Robert Stalnaker, "A Theory of Conditionals," *American Philosophical Quarterly*, monograph no. 2 (1968), pp. 98-112. See also the suggestions of Jonathan Bennett (and the people he cites) in "Counterfactuals and Temporal Direction," *The Philosophical Review* 93 January 1984, pp. 57-91.

there is a causal chain which produced (or sustains) S 's (noninferential) belief that p , and each link of it is a fact f which is both a reliable indicator as well as a conclusive indicator that p ; i.e.,

(1) at most times t , the following condition holds:

a member of S 's community is caused to believe that p^* at t by the fact that f^* holds at $t \rightarrow p^*$ at t ; and

(2) f holds $\rightarrow p$.²⁸

We arrived at this analysis by first identifying noninferential belief with nonaccidentally correct belief, and then working out the elements of the latter. Therefore our analysis will double as an account of when it is that we are not correct about our beliefs merely by luck.

2. Inferential Knowledge

It would be marvelous if the conditions for noninferential knowledge could simply be extended to inferential knowledge, so that the complete account of knowledge would already be before us (and the work behind us). Things are not so simple. I will show why we must complicate our theory, then complicate it.

Inferential beliefs, I assume, are ones we acquire by applying one or more inference rules. Certainly more enters into the causal history of such beliefs, but one link in a causal chain that produces an inferential belief will be a fact that fits into the following schema:

Person S has applied rule R so as to reach R 's recommendation that p be inferred.

Imagine what the world would be like if whenever we acquired an inferential belief through a causal chain each link of which is both a reliable and conclusive indicator, we would know that belief to be true. The fact that we have reached R 's recommendation is a link in any such chain; hence that fact must be both a reliable and conclusive indicator. This means that our reaching R 's recommendation must have the propensity in normal circumstances to accurately indicate that p , and that our present circumstances are not rigged. Hence our rule-governed beliefs must typically be acquired through the application of rules that recommend belief only if we have good evidence. Moreover, if we are in circumstances like

²⁸ This analysis replaces the one I offer in *The Epistemic Predicament*, op. cit.

those of the gemologist who was surrounded by diamond look-alikes, we must acquire our beliefs through rules that demand especially good evidence. This is because (the recommendation of) one rule will be more reliable than (the recommendation of) another primarily because the former requires better evidence than the latter before it will recommend belief.

The upshot is that we already have an account which usually demands that the quality of our evidence must exceed a certain minimum, and if our circumstances are ones in which it is difficult to get at the truth, our evidence must nonetheless be good enough for it to be the case that our circumstances are not rigged relative to our belief and its evidence. Moreover, our account can only be met by nonaccidentally correct beliefs. So why should we demand any more of a belief before we will count it as an instance of knowledge?

The problem is that some beliefs about which we are nonaccidentally correct would be *irrational*. To see why, consider the following crazy inference rule:

- R: For all S , if: S (believes that S) seems to see a table at some time, or S (believes that S) seems to see a dozen centaurs pulling the sleigh of a fat, jolly elf at some time, *infer*: There is a table in front of S at that time.

Sue Seesit, having seen a real table, then arrives at her belief that a table is in front of her by applying R. The fact that she has reached R's endorsement of her belief is both a reliable and a conclusive indicator that her belief is true. For the worlds near to the actual world in which she and those in her community reach R's recommendation are ones in which they seem to see a table. Rarely will there be worlds in which those in Seesit's community seem to see centaurs and elves; so rarely would R lead anyone astray.

Although Seesit's belief has as its source a causal chain each link of which is both a conclusive and a reliable indicator, so that she is nonaccidentally correct about her belief, she clearly fails to know that her belief is true. Beliefs acquired through rules like R are *irrational*. And irrational beliefs fall short of knowledge.

We can identify what is wrong with beliefs arrived at through rules such as R once we form a clear picture of what we expect out of the system of nondeductive inference rules whose application results in rational belief. I think that we can get at what we want by reminding ourselves that we intend to use this system to figure out our situation no matter what our situation might be. We do not want to use a system that we can rely on only if we are in a situation very much like the one in which we happen to

be. It is precisely this generality of dependableness that rules like R lack. Such rules are reliable only if a great deal is true of our situation, of the possible world we occupy. R itself is reliable only in a world in which very few centaurs and elves dwell. It is reliable only if our world has many of the features that it in fact has; in this sense its reliability is *world-bound*.

We want the reliability of our system of inductive rules to depend as little as possible on the way the world is; we want it to be as dependable as possible across as wide a range of the possible worlds we occupy as can be. Instead of the world-bound sort, then, we want our rules to have *inter-world* reliability. But our goal is not just to reduce the ranks of our false beliefs. First, we are not interested in avoiding false beliefs of just *any* sort. No one cares if they are wrong about the number of dust motes on Jupiter. And second, we also want to increase the ranks of our true beliefs, or rather those true beliefs relevant to our important interests. Moreover, this task of minimizing the ranks of those of our false beliefs that are relevant to our basic interests, while maximizing the ranks of our relevant true beliefs, we want our inductive rules to accomplish across as many of the possible worlds we occupy as can be.

I suggest that when the system of inductive rules through which we have acquired (or sustained) our beliefs has a goodly amount of interworld reliability in this interest-relative sense, those beliefs are rational. (Exactly *how much* is required cannot be specified.) Let me emphasize, however, that it is the entire system of rules with which we operate that must be interworld reliable, not a component of it. Presumably our logic will include rules that specify when a belief is *prima facie* justified, as well as ones which tell us when a *prima facie* justified belief is outweighed by conflicting evidence. A belief which is merely *prima facie* justified may not be rational, even if the rules governing *prima facie* justification are interworld reliable. It will *not* be if it is outweighed by overwhelmingly conflicting evidence according to the rules of our logic.²⁹ Even if it is endorsed by *all* of the rules of our logic it may still be irrational, however. Any logic, no matter how reasonable and reliable, can be crippled with additional rules, say ones which outweigh the good rules yet recommend absurd beliefs. If our logic's interworld reliability is undermined by such ridiculous rules, then the beliefs it yields will be irrational. *How* irrational these beliefs will be depends on just how interworld unreliable our logic is.

²⁹ In this way I believe I handle potential difficulties of the sort discussed by Lawrence Bonjour in "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy*, 5, 1980, ed. P. French, T. Uehling and H. Wettstein, pp. 53-73.

Having sketched an account of rational belief, we are now in a position to analyze inferential knowledge. We know of a noninferential belief that it is true if and only if it not only meets the conditions for noninferential knowledge, but is rational as well.³⁰

³⁰ A qualified version of Dretske's thesis that knowledge is information-caused belief is, I believe, entirely correct, namely that *noninferential* knowledge is information-caused belief. (To know that our inferential beliefs are true requires more than the fact that they share their informational content with their causes: they must also be justified.) The qualified thesis may be tested by working out the shape information must take in order to accommodate the thesis, then examining the resulting conception of information for plausibility.

Suppose that a noninferential belief is one that shares its informational content with one of its causes or sustainers. That is the qualified thesis. Then the fact that my belief and its source share their content must be necessary and sufficient for my (noninferentially) knowing that belief to be true. Hence an account of the informational content of a fact (or event) is ready to hand: a fact *f* carries the information that *p* (relative to person *S*) just in case *f*'s holding is both a reliable and conclusive indicator that *p* for *S*. And causal chains carry such and such information if and only if each of their individual links do.

I wish to thank Jonathan Bennett, Douglas Odegard, my colleagues at Trinity and an anonymous reviewer for helpful comment and criticisms. Part of this paper was presented at the 1985 Meeting of the Canadian Philosophical Association. My research was supported partly by a summer stipend from the National Endowment for the Humanities, and partly by a grant from Trinity University.