

Perception, Knowledge, and Belief

Selected Essays

FRED DRETSKE

Stanford University



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

© Fred Dretske 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United States of America

Typeface Bembo 10.5/13 pt. *System* DeskTopPro/UX [BV]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Dretske, Fred I.

Perception, knowledge, and belief : selected essays / Fred
Dretske.

p. cm. – (Cambridge studies in philosophy)

ISBN 0-521-77181-1. – ISBN 0-521-77742-9 (pbk.)

1. Knowledge, Theory of. 2. Philosophy of mind. I. Title.

II. Series.

BD161.D73 2000

121'.092—dc21

99-15973

CIP

ISBN 0 521 77181 1 hardback

ISBN 0 521 77742 9 paperback

Contents

<i>Preface</i>	<i>page ix</i>
----------------	----------------

Part One Knowledge

1	Conclusive Reasons	3
2	Epistemic Operators	30
3	The Pragmatic Dimension of Knowledge	48
4	The Epistemology of Belief	64
5	Two Conceptions of Knowledge: Rational vs. Reliable Belief	80

Part Two Perception and Experience

6	Simple Seeing	97
7	Conscious Experience	113
8	Differences That Make No Difference	138
9	The Mind's Awareness of Itself	158
10	What Good Is Consciousness?	178

Part Three Thought and Intentionality

11	Putting Information to Work	195
12	If You Can't Make One, You Don't Know How It Works	208
13	The Nature of Thought	227

14	Norms, History, and the Constitution of the Mental	242
15	Minds, Machines, and Money: What Really Explains Behavior	259
	<i>Index</i>	275

Conclusive Reasons

Conclusive reasons have a modal as well as an epistemic character. In having conclusive reasons to believe that P is the case one's epistemic credentials are such as to eliminate the possibility of mistake. This, at least, is how I propose to understand them for the remainder of this essay. Letting the symbol " $\langle \rangle$ " represent the appropriate modality (a yet to be clarified sense of *possibility*), I shall say, then, that R is a conclusive reason for P if and only if, given R, $\sim \langle \rangle \sim P$ (or, alternatively, $\sim \langle \rangle (R. \sim P)$). This interpretation allows us to say of any subject, S, who believes that P and who has conclusive reasons for believing that P, that, given these reasons, he *could not be wrong* about P or, given these reasons, *it is false that he might be mistaken* about P.

Suppose, then, that

- (1) S knows that P and he knows this on the basis (simply) of R entails
- (2) R would not be the case unless P were the case.¹

The latter formula expresses a connection between R and P that is strong enough, I submit, to permit us to say that if (2) is true, then R is a conclusive reason for P. For if (2) is true, we are entitled not only to deny that, given R, not-P is the case, but also that, given R, not-P *might* be the case. That is to say, (2) eliminates R and not-P as a possible

Reprinted From *Australasian Journal of Philosophy* 49, (1) (1971) 1–22, by permission of the publisher. An early version of this essay was read to the Philosophy Department at The University of Illinois, Chicago Circle. More recently it was read at The University of North Carolina's Colloquium in Philosophy, with Professor Robert Sleigh commenting.

¹ I shall be using "R" and "P" to replace a variety of related grammatical units. Depending on the sentential context, they sometimes serve as noun phrases, sometimes as full indicative sentences, sometimes for appropriate transformations of the indicative.

(joint) state of affairs and, when we are *given* R , it eliminates not- P as a possible state of affairs. This is so because (2) entails the falsity of

(3) Although R is the case P might not be the case.

If we express (3) as “Given R , $\langle \rangle \sim P$,” then (2) entails that it is false that, given R , $\langle \rangle \sim P$ that is equivalent to, given R , $\sim \langle \rangle \sim P$; and this is precisely the required feature of conclusive reasons given previously. Hence, when (2) is true, R is a conclusive reason for P .

What follows is an amplification of the preceding sketch – hence, an argument for the view that in those cases where knowledge that P rests on evidence, grounds, or reasons, when the question “How does S know?” can sensibly be asked and answered, the evidence, grounds, or reasons must be conclusive. Anything short of conclusive reasons, though it may provide one with justified true beliefs, fails to give the kind of support requisite to knowledge. I shall also urge that the possession of conclusive reasons to believe, properly qualified, is also a sufficient condition for knowledge.

1. KNOWING P ON THE BASIS OF R : THE CONNECTION BETWEEN (1) AND (2)

Suppose S , in order to assure himself that his child’s temperature is normal (no fever), places a thermometer in the child’s mouth, extracts it after several minutes, and observes a reading of 98.6°F. In remarking to the doctor that his child’s temperature is normal, S is asked how he knows. S responds, naturally enough, by saying, “I just took his temperature.” Let us assume, then, that we have an instantiation of (1):

(1a) S knows that his child’s temperature is normal, and he knows this on the basis of the (normal) reading on the thermometer (which he has just placed in the child’s mouth, etc.).

Can one consistently affirm (1a) and deny the corresponding instantiation of (2)?

(2a) The thermometer would not have read 98.6°F unless the child’s temperature was normal.

If it is not already obvious that one cannot consistently affirm (1a) and deny (2a), I think it can be made obvious by considering the kind of thing that would show (2a) to be false. For example, if Charles, familiar with the particular thermometer in question, should say, “Oh, I know

that thermometer; it is fairly accurate for temperatures below 98° but it sticks at 98.6 for almost any higher temperature,” we have been given solid grounds for rejecting (2a). Simultaneously, however, we have been given solid grounds for rejecting (1a). If it is *that* kind of thermometer, then if S’s only basis for thinking his child’s temperature normal is a 98.6 reading on it, then he does not *know* that his child’s temperature is normal. It *might* be normal, of course, but if S knows that it is, he must have more to go on than the reading on this (defective) thermometer.

Other attempts to show (2a) to be false have the same effect; they immediately undermine R (the reading on the thermometer) as an adequate basis for someone’s knowing that the child’s temperature is normal (P). For in rejecting (2a) we reject the thermometer as a reliable device for discovering whether a person’s temperature is normal, and knowledge is not acquired by relying on what is unreliable in precisely those respects in which we rely on it.

We frequently purport to know things on the basis of testimony. James has a large stamp collection and, after giving a detailed description of it, invites S to his home to see it. S declines, but he later refers to James’s stamp collection in conversation. It is easy enough to imagine circumstances in which it would be natural for someone to ask S how he knew, what reasons he had for thinking that James had such a collection. And it is just as easy to imagine S, in response to such a query, referring to his conversation with James. Let us assume, then, that

(1b) S knows that James has a stamp collection, and he knows this on the basis of James’s description (and invitation)

is true. I am not now concerned to argue that one can know something of this sort in this way; the question is, rather, whether (2b) must be true *if* (1b) is true.

(2b) James would not have said he had a stamp collection, described it in such detail, and issued an invitation unless he had a stamp collection.

If James is the sort of fellow about which (2b) cannot be truly asserted, then he is not the sort of fellow who can be trusted on such matters as this. If James is the sort of fellow who sometimes says such things as a joke, who would (or might) concoct such elaborate stories for his own amusement (or whatever), who would (or might) whimsically issue an invitation of this sort under totally false pretexts, then, despite the fact that he is (by hypothesis) telling the truth on this occasion, his testimony is hardly the sort on which one can rest a claim to know. In denying

(2b) one is conceding that James would, or might, have said what he did without possessing a stamp collection, and in the light of this concession one cannot go on to insist that, nonetheless, S *knows* he has a stamp collection on the basis, simply, of what James said.

Gilbert Harman contrasts two cases, *the lottery case* and *the testimony case*.² Although S, say, has only one among thousands of tickets in a lottery and, hence, has an extremely slender chance of winning, we naturally reject the idea that S could know that he was going to lose on the basis of a correct probability estimate (well over 99.9%) of his losing. Even if S correctly predicts that he is going to lose, we would deny that he knew he was going to lose if the *only* basis he had for this belief was the fact that his chances of winning were so slight.³ Harman compares this case with the situation in which we often seem prepared to say that S knows that P when he is told that P is the case by some other person (testimony case). Although probability estimates are not altogether appropriate here, we do know that people sometimes lie, sometimes they are honestly mistaken, and so on. There always seems to be a chance that what a person tells us is not the case, however sincere or credible he may appear, and the order of magnitude of this chance seems to be comparable to the chance that we might win in some appropriate lottery situation. Why, then, are we prepared to say that we know in the one case but not in the other? Harman has some revealing remarks to make about these cases, but I mention them only to bring out their relevance to the present discussion. For I think this contrast strengthens the view that (2) is normally accepted as a necessary consequence of (1), that when we are unwilling to endorse the corresponding instantiation of (2) we are unwilling to talk of anyone knowing that P is the case on the basis of the evidence expressed by R. In many testimony situations we are, I believe, willing to affirm (2): the person would not have said it unless it was so. In the lottery case, however, the connection between R and P expressed by this subjunc-

2 "Knowledge, Inference, and Explanation," in *American Philosophical Quarterly*, July 1968.

3 Of course S may have said "I know I am going to lose" and he may say now, after he has lost, "I knew I was going to lose," but these expressions are normally accepted without epistemological quibbling because they are taken as little more than expressions of resignation or despair. With this use of the verb "to know," one can know one is going to lose and *still* spend a dollar for a ticket and a chance at the prize—a fact about human beings that is puzzling if they believe they know (in any epistemologically relevant sense) that they are going to lose.

tive conditional fails to be realized, and it fails no matter how great the probabilities become. Adjusting the wording of (2) to suit the example in question⁴ we have

- (2c) If S were going to win the lottery, his chances of winning would not be $1/m$ (m being the number of tickets sold).

Whatever (finite) value we give to m , we know this is false since someone whose chances of winning are $1/m$ *will* win, and since there is nothing special about S that would require him to have a better chance of winning than anyone else in order to win, we reject (2c) as false. Hence, we reject the idea that S can know he is going to lose on the basis of the fact that his chances of losing are $(m - 1)/m$.

Alvin Goldman, in developing a causal account of knowledge, constructs a situation in which S is said to know that a nearby mountain (I will call it M) erupted many years ago. He knows this on the basis of the presence of solidified lava throughout the countryside surrounding the mountain.⁵ According to Goldman, a necessary condition for S's knowing that M erupted many years ago on the basis of the present existence and distribution of the lava is that there be a causal connection between the eruption of the mountain and the present existence and distribution of the lava. I do not wish to dispute this claim at the moment since the view I am advancing is even stronger: viz. that a necessary condition for S to know that M erupted on this basis is that

4 The wording of (2) will sometimes have to be adjusted to suit the particular instantiation in question. The chief factors determining this adjustment are the relative temporal locations of R, P, and the time of utterance and also the causal connections, if any, that are believed to hold between R and P. The particular wording I have given (2) is most appropriate when P is some state of affairs antecedent to (or contemporaneous with) both R and the time of utterance. This, of course, is the result of the fact that (2) is most often used when P is some state of affairs causally responsible for the present condition R. When P is a future state we might express (2) as: R would not be the case unless P were going to happen. For example, he would not have registered unless he were going to vote. I do not wish to preclude the possibility of knowing that something *will* occur on the basis of present evidence by restricting the wording of (2). The difficulty, of course, is that when P is some future state, the subjunctive relating it to R generally becomes somewhat questionable. We prefer to say, in our more cautious moods, that if he were not *planning* to vote, he would not have registered (acknowledging, thereby, the fact that contingencies may interfere with the execution of his plans). But in the same cautious moods we prefer to say, not that we know he is going to vote (because he registered), but that we know he plans or intends to vote.

5 "A Causal Theory of Knowing," *Journal of Philosophy*, June 22, 1967, p. 361.

(2d) The lava would not be here, and distributed in this manner, unless M erupted

is true. (2d) is a stronger claim than that the eruption of M is causally connected with the present existence and distribution of the lava. (2d) requires, in addition, that M's eruption be necessary for the present state of affairs. To illustrate, consider the following embellishment on Goldman's example. Not far from M is another mountain, N. The geology of the area is such that at the point in time at which M erupted something, so to speak, was bound to give; if M had not erupted, N would have. Furthermore, the location of N is such that if it, rather than M, had erupted, the present distribution of lava would have been, in all respects relevant to S's taking it as a reason for believing M erupted, the same. In such circumstances Goldman's necessary condition is satisfied, but mine is not. (2d) is false; it is false that the lava would not be here, and distributed in this fashion, unless M had erupted. For if, contrary to the hypothesis, M had not erupted, N would have, leaving the very same (relevant) traces.

In such circumstances I do not think we could say that S knew that M erupted on the basis of the present existence and distribution of lava. For, by hypothesis, *this* state of affairs would have obtained whether M erupted or not and, hence, there is nothing about this state of affairs that favors one hypothesis (M erupted) over a competing hypothesis (N erupted). S is still correct in supposing that M did erupt, still correct in supposing that it was M's eruption that is causally responsible for the present existence and distribution of lava, but he does not know it was M that erupted – not unless he has some additional grounds that preclude N. If he has such additional grounds, call them Q, then we can say that he knows that M erupted and he knows this on the basis of *R* and *Q*. In this case, however, the corresponding instantiation of (2) is also satisfied: *R* and *Q* would not be the case unless M erupted. As things stand, the most that S could know, on the basis simply of the present existence and distribution of lava, is that *either M or N erupted*. (2) permits us to say this much, and no more, about what can be known on the basis of lava flow.

The case becomes even clearer if we exploit another of Harman's examples.⁶ Harold has a ticket in a lottery. The odds against his winning

6 In "Knowledge, Inference, and Explanation," pp. 168–169. I have adapted the example somewhat.

are 10,000 to 1. The prize, call it an X, is something that Harold does not now have nor could he reasonably expect to obtain one by means other than winning the lottery. Enter a philanthropic gentleman, Rockaford, who decides to give Harold an X if (as seems likely) he should fail to win one in the lottery. Things go as one would expect; Harold holds a losing ticket. Rockaford keeps his word and gives Harold an X. S, familiar with the preceding circumstances but unaware of whether Harold won or lost in the lottery, finds Harold with his newly acquired X. S infers that Harold received his X from Rockaford. He concludes this because he knows that the only other way Harold might have acquired an X is by winning the lottery and the odds against that happening are enormous. The following conditions are satisfied: (a) Harold received his X from Rockaford; (b) S believes that Harold received his X from Rockaford; (c) S is warranted in believing this since the chances of his having received it in any other way are negligible; (d) Rockaford's generous gift of an X to Harold is the (causal?) explanation of Harold's present possession of an X; and, finally (e) S correctly reconstructs (to use Goldman's language) the causal chain of events that brought about Harold's possession of an X. Yet, why does it seem clear (at least to myself – and apparently to Harman) that S does *not* know that Rockaford gave Harold his X. Because

(2e) Harold would not have an X unless Rockaford gave him one

is plainly false.⁷ If Rockaford had not given him an X, it would have been because Harold already possessed one as a winner in the lottery. Hence, Harold would possess an X even if Rockaford had not given him one. It is not true that R would not be the case *unless* P; hence, not true that S knows that P on the basis of R.⁸

(2), therefore, expresses something stronger than a causal relationship be-

7 There is a way of reading (2e) that makes it sound true – viz., if we illicitly smuggle in the fact that Harold *has lost* the lottery. That is, (2e') "Harold, having lost the lottery, would not have an X unless Rockaford had given him one" is true. But this version of (2) makes R, the reason S has for believing Rockaford gave him an X, not only Harold's possession of an X but also *his having lost the lottery*. This, by hypothesis, is not part of S's reason; hence, not properly included in (2e). (2e) must be read in something like the following fashion: Harold would not have an X, whatever the outcome of the lottery, unless Rockaford had given him one. With this reading it is clearly false.

8 It is difficult to say whether this is a counterexample to Goldman's analysis. I think it satisfies all the conditions he catalogs as sufficient for knowledge, but this depends on how strictly Goldman intends the condition that S must be warranted in inferring that P is the case from R.

tween R and P. It should be pointed out, however, that it expresses something that is, in certain important respects, weaker than a universal association between states or conditions similar to R and states or conditions similar to P. When “R” and “P” are expressions that stand for particular conditions or states of affairs, as will often be the case when we know one thing on the basis of something else, (2) expresses a connection between more *determinate* states of affairs than those described by talking about states similar to R and P. If someone remarks, midway through a poker hand, that if his neighbor had not folded (dropped from the game) he (the speaker) would have been dealt a royal flush, he is obviously not maintaining that *whenever* his neighbor remains in the game, he (the speaker) is rewarded with a royal flush. Rather, he is talking about *this* hand (already holding four cards to the royal flush), *this particular* distribution of cards in the remainder of the deck, *this particular* seating arrangement, and so on. He is not saying that his neighbor’s remaining in the game is, quite generally, sufficient for his receipt of a royal flush. Rather, he is saying that *in the particular circumstances which in fact prevailed on this occasion*, circumstances that include such things as card distribution, arrangement of players, and so on, an occurrence of the first sort (neighbor remains in game) will invariably be followed by one of the second sort (his receipt of a royal flush). One cannot falsify his claim by showing that he would not have received a royal flush, despite his neighbor’s remaining in the game, if the card distribution in the deck had been different from what it in fact was. For his claim was a claim about the inevitable sequence of events *with that distribution of cards*.

Statements such as (2), then, even when R and P are expressions for particular states of affairs, express a general uniformity, but this general uniformity is not that whenever a state similar to R is the case, then a state similar to P will also be (or have been) the case. The uniformity in question concerns the relationship between states similar to R and P *under a fixed set of circumstances*. Whenever (a state such as) R *in circumstances C* then (a state such as) P where the circumstances C are defined in terms of those circumstances that actually prevail on the occasion of R and P. But does C include *all* the circumstances that prevail on the occasion in question or only *some* of these? Clearly not all the circumstances since this would trivialize every subjunctive conditional of this sort. Even if we restrict C to only those circumstances logically independent of R and P we still obtain a trivialization. For, to use Goldman’s mountain example (as embellished), C would still include the fact that N did not erupt (since this is logically independent of both R and P),

and this is obviously one of the circumstances *not* held fixed when we say that the lava would not be here unless M erupted. For in asserting his subjunctive we mean to be asserting something that would be *false* in the situation described (N would have erupted if M had not), whereas if we hold this circumstance (N did not erupt) fixed, the uniformity between the presence of lava and M's eruption would *hold*.

I think that our examples, not to mention an extensive literature on the subject,⁹ point the way to a proper interpretation of C. The circumstances that are assumed constant, that are tacitly held fixed, in conditionals such as (2), are those circumstances prevailing on the occasion in question (the occasion on and between which the particular states R and P obtain) that are logically and causally independent of the state of affairs expressed by P.¹⁰ When we have a statement in the subjunctive that (unlike (2)) is counterfactual (the antecedent gives expression to a state of affairs that does or did not obtain), then C includes those circumstances prevailing on the occasion that are logically and causally independent of the state of affairs (or lack of such state) expressed by the *antecedent* of the conditional. In our poker game, for example, we can say that S's statement (I would have got a royal flush if my neighbor had stayed in the game) *fixes* that set of circumstances that are logically and causally independent of his neighbor's staying in the game (i.e., the antecedent since the statement is counterfactual). Hence, if there is another player in the game (whose presence or absence affects the cards dealt to S) who would have dropped out if S's neighbor had not dropped

9 I am not proposing a solution to the problem to which Nelson Goodman (*Fact, Fiction and Forecast*, Chapter I), Roderick Chisholm ("The Contrary-to-Fact Conditional," *Mind*, 55, 1946), and others addressed themselves in trying to specify the "relevant conditions" associated with counterfactuals. I shall use the notion of "causality" in my treatment, a device that both Goodman and Chisholm would regard as question-begging. I am not, however, attempting to offer an extensionalist analysis of the subjunctive conditional; I am merely trying to get clear in what way such conditionals are stronger than the statement of a causal relationship between R and P and yet (in one sense) weaker than a statement of the universal association between states similar to R and P.

10 This characterization of the circumstances "C" has interesting and, I believe, significant repercussions for subjunctives having the form of (2) in which R expresses some present (or past) state of affairs and P expresses some future state of affairs. Although I lack the space to discuss the point here, I believe an important asymmetry is generated by a shift in the relative temporal locations of R and P. I also believe, however, that this asymmetry is faithfully reflected in the difference between knowing what *will* happen on the basis of present data and knowing what *did* happen on the basis of present data. In other words, I feel that an asymmetry in (2), arising from a shift in the relative temporal locations of R and P, helps one to understand the difference we all feel between knowing, on the one hand, what did happen or is happening and, on the other hand, what will happen.

out, then this person's remaining in the game is *not* held fixed, not included in C, because it is causally connected to the state of affairs expressed by the antecedent in S's statement. Therefore, we can show S's statement to be false if we can show that such a circumstance prevailed, and it is along these lines that one would surely argue in attempting to show S that he was wrong, wrong in saying that he would have received a royal flush if his neighbor had stayed in the game.

On the other hand, one cannot show that S's statement is false by showing that, were the cards differently arranged in the remainder of the deck, he would not have received his royal flush; for the arrangement of cards in the remainder of the deck (unlike the presence or absence of our other player) is (presumably) independent of S's neighbor's departure from the game. Hence, it is one of the conditions held fixed, included in C, by S's statement, and we are not allowed to consider alterations in it in assessing the general implication of S's statement.

Or consider our original thermometer example. Recall, the statement in question was: "The thermometer would not have read 98.6° unless the child's temperature was normal." Suppose someone responds, "Oh, it would have (or might have) if the thermometer was broken." It is important to understand that one can grant the truth of this response without abandoning the original assertion; for the original assertion had, as its general implication, *not* a statement expressing a uniform relationship between states of affairs similar to the child's temperature (normal body temperature) and states of affairs similar to the thermometer reading (a reading of 98.6), but, rather, a uniformity between such states *under a fixed set of circumstances*. And, if I am right, this fixed set of circumstances includes *the actual state of the thermometer* (defective or accurate); it is one of those circumstances prevailing on the occasion in question that is causally and logically independent of the child's temperature. Hence, this circumstance cannot be allowed to vary as it is in the preceding response by the words "if the thermometer was broken." To determine the truth of the original assertion, we must suppose that the thermometer is accurate (or defective) *whatever the actual condition is*. If, therefore, the thermometer was not broken or otherwise defective on that occasion, then the suggestion that it would (or might) have read the same despite a feverish child if it were broken or defective is, although quite true, irrelevant to the truth of the statement: "It would not have read 98.6 unless the child's temperature was normal."

One final important feature of (2). I have said that, generally speaking,

the placeholders “R” and “P” represent expressions designating *specific* states of affairs or conditions. When this is so, (2) still has a general implication, but the general implication, expressing a uniform relationship between states of affairs similar to R and P, *has its scope restricted* to situations in which the circumstances C (as specified earlier) obtain. Since we are talking about specific states of affairs in most instantiations of (2), it becomes extremely important to observe the sorts of referring expressions embodied within both “R” and “P.” For example, when I say, “John would not have said it was raining unless it was raining,” I am talking about *John* and about a *particular utterance* of his. *Someone else* might have said this without its being true; John may have said *something else* without its being true. Nonetheless, *John* would not have said *it was raining* unless it was. An incurable liar about most things, John has a pathological devotion to accuracy on matters concerning the weather. In such a case, although John is, admittedly, a most unreliable informant on most matters, we can say that he would not have said it was raining unless it was so. This is only to say that the referring expressions to be found in “R” and “P” *help to define the scope* of the implied generalization. Recall, the implied generalization was about states of affairs *similar* to (the particular states) R and P. Similar in what respect? The sorts of referring expressions to be found within “R” and “P” help us to answer this question. In the case of John, the general implication involved not *a person’s saying something* (under circumstances C), not *John’s saying something* (under circumstances C), but John’s saying something about the weather (under circumstances C).

2. THE POSSIBILITY OF MISTAKE: THE CONNECTION BETWEEN (2) AND (3).

Taking a cue from the fact that (2) expresses some form of necessary relationship between R and P, I have (in my opening remarks) expressed (2) as: Given R, $\sim \langle \rangle \sim P$ (or, alternatively, $\sim \langle \rangle (R. \sim P)$). I think the full justification for expressing (2) in this fashion lies in the fact that (2) and (3) are contradictories, and since (3) may be rendered as

(3) Given R, $\langle \rangle \sim P$ (or, alternatively, $\langle \rangle (R. \sim P)$)

(2) may be represented as “Given R, $\sim \langle \rangle \sim P$.”

To see that this is so, it should be noticed that in denying the connection between R and P expressed by (2) we do not commit ourselves to anything as strong as

(4) R would be the case even though not-P were the case.

(4) is the *contrary* of (2), not its contradictory, since both (2) and (4) may turn out false.¹¹ For example, suppose S asserts,

(2g) I would have won the lottery if I had bought two tickets (instead of only one).

We may deny the truth of this contention without committing ourselves to the truth of

(4g) You would have lost even if you had bought two tickets.

All that is intended in denying (2g) is that the purchase of two tickets is connected with winning in the alleged manner, that the purchase of two tickets would have assured him of winning. Its failing to be connected in the manner alleged is, however, quite consistent with *his winning* with two tickets. What we commit ourselves to in denying (2g) is

(3g) You *might* have lost even with two tickets.

(3g) asserts what (2g) denies; viz. that even with two tickets it is *still* a matter of chance; the possibility of losing is *not* eliminated by holding two tickets instead of one.

As a matter of common practice, of course, we often employ something similar to (4) in denying (2). This is understandable enough since the truth of (4) does entail the falsity of (2). The point I am driving at, however, is that we need not affirm anything as strong as (4) in denying (2); all we are required to affirm is that R and not-P *might* both be the case or that, even though R is given, P *might not* be the case. That is to say, the proper expression for the negation of (2) is (3); and if we understand (3) as affirming “Given R, $\langle \rangle \sim P$ ” (alternatively $\langle \rangle(R. \sim P)$), then we are justified in representing (2) as “Given R, $\sim \langle \rangle \sim P$ ” (alternatively, $\sim \langle \rangle(R. \sim P)$). If someone says, “James would not have come without an invitation,” we can deny this without supposing that James would have come without an invitation. For suppose we know that if James had not received an invitation, he would have flipped a coin to decide whether to go or not. In such a case, it is not true that he would not have come without an invitation, but neither is it true

11 A point that Goodman acknowledges in a footnote: “Literally a semifactual and the corresponding counterfactual are not contradictories but contraries, and both may be false.” *Fact, Fiction and Forecast*, p. 32, note 2.

that he would have come without an invitation. The fact is that *he might have come* (depending on the outcome of the toss) without an invitation.

Before proceeding there is an important ambiguity that must be eliminated. There is a use of the word “might” (and related modal terms) that gives expression to the speaker’s epistemic state in relation to some situation. It is a use of the word “might” that is naturally followed by a parenthetical “for all I know.” For example, in saying, “He might win the election,” the speaker is most naturally understood as expressing the fact that he does not know whether he will win the election or not, that he has no (or there are no) very convincing grounds for supposing he will lose. This use of the word can properly be deployed even when the state of affairs in question is *physically* or *logically* impossible. S, for instance, may concede the premises of our valid argument but, nonetheless, ignorant of its validity, insist that the conclusion might (for all he knows) be false even though the premises are true.

Contrasted with this epistemic use of the modal terms is what we might call an *objective* sense of these terms, a sense of the term “could” (for example) in which if R entails P then, independent of what S knows about the logical relationship between R and P, it is false to say that R and not-P could (or might) both be the case. Moreover, if we accept the results of modern physics, then in this objective sense of the term S’s statement that there are objects that *can* travel faster than the speed of light is false, and it is false even though, *for all he knows*, there are objects that can. In this objective sense one is making a remark about the various possibilities for the occurrence, or joint occurrence, of events or the coexistence of states of affairs, and in this sense ignorance of what is the case is no guarantee that one’s statements about what *might* or *could* be the case are true. The possibilities must actually be as one alleges. When S (knowing that James had an invitation to come) asserts that James, being the sort of fellow he is, might (even) have come without an invitation, he is making a remark *about James* and what James is inclined to do or capable of doing. He is obviously not registering some fact about his ignorance of whether or not James possesses an invitation.

The modal term appearing in (3) is meant to be understood in the objective sense. (3) is meant to be a statement about the possibilities for the joint realization of two states of affairs (R and not-P) independent of what the speaker happens to know about the actual realization of P (R being given). Drawing from our discussion in the preceding section, we can say that if (2) is true, if R would not be the case unless P were the case, then *in these circumstances* (specified earlier) P is a state of affairs

that is necessary to the realization of R. Hence, in these circumstances it is false to say that R and not-P might both be the case or that, given R, not-P might be the case, and it is false whether or not anyone appreciates the fact that it is false.

We have here a more or less *particularized* impossibility; (3), as well as (2), is tied to the circumstances, C, specified earlier. Nothing else but (a state such as) P could have brought about (a state such as) R *in these circumstances*. Often, of course, our statements about what could be the case, about what is possible, are broader in scope. They do not restrict themselves to the particular set of circumstances prevailing on some specific occasion. They are statements to the effect that, whatever the circumstances on this occasion happened to be, there are (nonetheless) circumstances in which a state relevantly similar to R is, or might easily be, brought about without a state relevantly similar to P. S may admit, for example, that the thermometer would not have read 98.6 unless the child's temperature was normal – hence, concede that it would be false to say that in these circumstances (given the thermometer reading) the child might (nonetheless) have a fever. Yet, he may go on to insist that one *can* get a normal reading on a thermometer with a feverish child. One can do so when one has a defective thermometer, when the child is sucking an ice cube, and so on. That is, one *can* have R without P in *other* circumstances. These “general” possibilities are, however, quite consistent with the “particularized” impossibilities expressed by (3). Most genuine impossibilities can be made possible by enlarging the frame of reference, by relaxing the conditions tacitly taken as fixed in the original statement of impossibility. I can't swim the English Channel; this despite the fact that I could *if* I had trained since a boy, been endowed with the requisite endurance, and so on.

If I may summarize the argument up to this point in preparation for the following section: (1) entails (2), if S knows that P, and he knows this on the basis (simply) of R, then R would not be the case unless P were the case. Furthermore, (2) gives expression to a connection between R and P that permits us to say that when R is given, in the kinds of circumstances that actually prevailed on the occasion in question, the possibility of not-P is eliminated. The sense of the word “possible” that is operative here is, I submit, the *same* sense of this word that is operating in our strongest statements about what is physically possible; the difference between the possibility expressed in (3) and other, apparently stronger, statements of what is possible and impossible is simply a shift in the set of circumstances that is taken as fixed. The impossibility

expressed by (3) is an impossibility relative to those circumstances, C, held fixed in (2), and it is this fact that makes (3) the contradictory of (2) and, hence, that makes its negation a logical consequence of (1).

3. CONCLUSIVE REASONS

Let us call R a conclusive reason for P if and only if R would not be the case unless P were the case.¹² This makes logically conclusive reasons (LCR) a subclass of conclusive reasons. Conclusive reasons depend, simply, on the truth of (2); LCR require that the truth of (2) be demonstrable on purely logical and definitional grounds. When the conditional is true, but not logically true, we can speak of the conclusive reasons as being *empirically conclusive* (ECR).

Of course, R may *be* a conclusive reason for believing P without anyone believing P, much less having R as his reason for believing. I shall say, therefore, that S *has conclusive reasons*, R, for believing P if and only if:

- (A) R is a conclusive reason for P (i.e., (2) is true),
- (B) S believes, without doubt, reservation, or question, that P is the case and he believes this on the basis of R,
- (C) (i) S knows that R is the case or
 - (ii) R is some experiential state of S (about which it may not make sense to suppose that S *knows* that R is the case; at least it no longer makes much sense to ask *how* he knows).

With only minor embellishments, to be mentioned in a moment, I believe that S's having conclusive reasons for believing P is *both* a necessary and a sufficient condition for his knowing that P is the case. The appearance of the word "know" in this characterization (in (Ci)) does not render it circular as a characterization of knowledge since it can be eliminated by a recursive application of the three conditions until (Cii) is reached.

If S has conclusive reasons for believing P, then it is *false* to say that, given these grounds for belief, and the circumstances in which these grounds served as the basis for his belief, S *might be mistaken about P*. Having conclusive reasons, as I have just defined it, not only implies that P is the case, it not only implies that S believes that P is the case,

12 Recall footnote 5 concerning the particular wording of (2); I intend those remarks to apply to this definition of "conclusive reasons."

but it also implies that, in the circumstances in which he came to believe that P, his basis for believing that P was sufficiently secure to eliminate the possibility of his being mistaken about P. This goes a long way toward capturing everything that philosophers have traditionally required of knowledge. Indeed, in certain respects it goes beyond it in requiring a stronger connection between one's reasons or grounds and what one believes (on the basis of these reasons or grounds) than has normally been demanded by those wishing to preserve our ordinary knowledge claims from skeptical criticism.¹³ Since, however, I have already argued for the necessity of (A), for R's being a conclusive reason for P, I shall concentrate on the question of whether or not S's having a conclusive reason for believing P is sufficient for his knowing that P.

Several preliminary points must be mentioned briefly. It may be thought that in arguing for the impossibility of mistake as a necessary condition for knowing that P, I have been wasting my time. It may be thought that if S knows that P, then P *cannot* be false since S's knowing that P entails P; hence, S *cannot* be mistaken in believing that P. In answer to this objection it should be pointed out that the impossibility of mistake that I have been talking about is an impossibility that arises in virtue of a *special connection* between S's reasons, R, and what he consequently believes, P. It is not the trivial impossibility of being wrong about something that (by hypothesis) you know. When philosophers

13 It is this stronger connection that blocks the sort of counterexample that can be generated to justified-true-belief analyses of knowledge. Gettier's (and Lehrer's) examples, for instance, are directed at those analyses that construe knowledge in terms of a *degree* of justification that is compatible with being justified in believing something false (both Gettier and Lehrer mention this feature at the beginning of their discussion). The counterexamples are then constructed by allowing S to believe that P (which is false) *with the appropriate degree of justification*, letting P entail Q (which is true), and letting S believe that Q on the basis of its logical relationship to P. We have, then, a case where S truly believes that Q *with the appropriate degree of justification* (this degree of justification is allegedly preserved through the entailment between P and Q), but a case where S does *not know* that Q (since his means of arriving at it were so clearly defective). On the present analysis, of course, the required connection between S's evidence and P is strong enough to *preclude* P's being false. One cannot have *conclusive* reasons for believing something that is false. Hence, this sort of counter-example cannot be generated. Part of the motivation for the present analysis is the conviction (supported by Gettier-like examples) that knowledge, if it embodies an evidential relation at all, must embody a strong enough one to eliminate the *possibility* of mistake. See Edmund Gettier's "Is Justified True Belief Knowledge?" *Analysis*, 23.6, June 1963, and Keith Lehrer, "Knowledge, Truth and Evidence," *Analysis*, 25.5, April 1965. I should also mention here that these same sorts of considerations seemed to move Brian Skyrms toward a similar analysis; see especially pp. 385–386 in his "The Explication of 'X knows That P'," *The Journal of Philosophy*, June 22, 1967.

concern themselves with the possibility of mistake in putative cases of knowledge, they are not concerned with the possibility of mistake that is trivially avoidable by saying that *if* you do know that P, then you cannot be mistaken about P. They are concerned, rather, with that possibility as it exists in relation to one's evidence or grounds for believing P, and *that* is the possibility with which (2) is concerned.

The point may be put in another way. Both

- I. R would not be the case unless P were the case
R is the case

and

- II. $R \supset P$ (when it is *not* true that R would not be the case unless P)
R

constitute conclusive grounds, logically conclusive grounds, for believing P. Neither set of premises would be true unless P were true, and this fact is in both cases demonstrable on purely logical and definitional grounds. But the significant difference between I and II is that in I, but *not* in II, the second premise *alone* turns out to be a *conclusive* reason (ECR). If we were searching for conclusive reasons to believe P, then in the second case we would require as our reasons *both premises*, and this would require that we knew that both premises were true (see clause (C) in having conclusive reasons). In case I, however, the second premise alone is a conclusive reason and, hence, to have conclusive reasons it is required *only that we know that R is the case*. We need not (as in case II) know that the first premise is true. All that is required in case I for R alone to be a conclusive reason is that the first premise be true; there is nothing that requires S to know that the first premise is true in order to have R as his conclusive reason for believing P. For if the first premise is true (regardless of whether S knows it is true or not), then (3) is false; hence, the possibility of S's being mistaken about P has been successfully avoided – and it has been successfully avoided *whether or not S knows it has been avoided*.

In speaking of conclusive reasons I do not wish to suggest that in having R as a conclusive reason S must be in a position to *give* R as his reason. R may simply be a certain experience that S has undergone and, having undergone this experience, come to the belief that P was the case on the basis of (as a result of) this experience. He may find it difficult, or impossible, to give verbal expression to R. He may have forgotten R. Or R may consist in something's looking a particular way

to him that he finds difficult to describe. Still, if the way the thing looks to S is such that it would not look that way unless it had the property Q, then its looking that way to S is a conclusive reason for S's believing that it has the property Q; and if S believes that it is Q *on this basis*, then he has, in the way the thing looks to him, a conclusive reason for believing it Q.

Also, there are a number of things that people commonly profess to know (Sacramento is the capital of California, the earth is roughly spherical) for which there is no definite piece of evidence, no single state of affairs or easily specifiable set of such states, that even approximates a conclusive reason. In such cases, although we can cite no single piece of data that is clinching and, hence, are at a loss for conclusive reasons when asked to give reasons (or when asked "How do you know?") we, nonetheless, often enough have conclusive reasons in a vast spectrum of experiences that are too diverse to admit of convenient citation. Countless experiences converge, so to speak, on the truth of a given proposition, and this variety of experience may be such that although one *may* have had any *one* of these experiences without the proposition in question being true, one *would not* have had *all* of them unless what one consequently believes was true. The fallibility of source A and the fallibility of source B does not automatically entail that when A and B agree about P's being the case, that, nonetheless, P might still be false. For it may be that A and B *would not* both have indicated that P was the case unless P was the case, although neither A nor B, taken by themselves, provide conclusive reasons for P. For example, although any single newspaper account may be in error on a particular point, several independent versions (wire services, of course, tend to eliminate this independence) may be enough to say that we know that something is so *on the basis of the newspaper accounts*. All of them would not have been in such close agreement unless their account was substantially correct.¹⁴

Finally, I do not wish to suggest by my use of the word "reason" that

14 The fact that all newspapers sometimes print things that are false does not mean that we cannot know that something is true on the basis of a single newspaper account. The relevant question to ask (as in the case of a person's testimony – see Section 1) is not whether *newspapers* sometimes print false stories, not even whether *this newspaper* sometimes prints false stories, but whether *this newspaper* would have printed *this story* if it were not true. The *Midville Weekly Gazette's* story about dope addiction on the campus may not correspond with the facts, but would *The Times* have printed this story about the president's visit to Moscow if it were not true?

when S has conclusive reasons for believing P, S has *reasoned* his way to the conclusion that P is the case from premises involving R or that S has consciously used R as a reason in arriving at the belief that P. I am inclined to think (but I shall not now argue it) that when one knows that P, on whatever basis this might be, little or no reasoning is involved. I would prefer to describe it as follows: sometimes a person's conviction that P is the case can be traced to a state of affairs (or cluster of situations) that satisfies the three conditions defining the possession of conclusive reasons. When it can be so traced, then he knows; when it cannot be so traced, then we say he does not know, although he may be right about P's being the case. Of course, his belief may be *traceable* to such a source without our being able to trace it. In such a case we are mistaken in saying that he does not know.

Turning now to the question of whether having conclusive reasons to believe, as defined by (A)–(C), constitutes a sufficient condition for knowledge, I shall mention and briefly respond to what I consider to be the most serious objections to this proposal.

There is, first, a tendency to conflate knowing that P with knowing that one knows that P. If this is done then conditions (A)–(C) will immediately appear insufficient since they do not describe S as knowing or having any basis for believing that R, his basis for believing P, constitutes an adequate basis, much less a conclusive basis, for believing P. Even if one does not go this far, there is still a tendency to say that if S knows that P, then S must at least *believe* that he knows that P is the case. If one adopts this view then, once again, conditions (A)–(C) appear inadequate since they do not describe (nor do they entail) that S believes he knows that P is the case. I see no reason, however, to accept either of these claims. We naturally expect of one who knows that P that he believe that he knows, just as we expect of someone who is riding a bicycle that he believe he is riding one, but in neither case is the belief a *necessary* accompaniment. The confusion is partially fostered, I believe, by a failure to distinguish between what is implied in knowing that P and what is implied (in some sense) by someone's *saying* he knows that P. Consider, however, cases in which we freely ascribe knowledge to agents in which it seems quite implausible to assign the level of conceptual sophistication requisite to their believing something about knowledge, believing something about their epistemic relation to the state of affairs in question. A dog may know that his master is in the room, and I (at least) want to say that he can know this in a straightforward sense

without (necessarily) possessing the conceptual understanding that seems to be required to say of him that he believes he knows this, believes that he has good reasons, or believes *anything* about his epistemic relation to the fact that his master is in the room.¹⁵ Yet, it seems perfectly natural to say of a dog that he knows (sometimes) when his master is in the room. And this is not, let me add, simply a matter of the dog's being right about something. For if we knew that the dog thought his master was in the room on the basis, say, of certain sounds and smells (sounds and smells that were, generally speaking, a reliable sign of his master) when these sounds and smells were totally unrelated to his master's presence in the room *on this occasion*, we would not say the dog knew even if he happened to be right about his master's being in the room. Imagine a blind dog's being "taken in" by a thief in his master's clothing while the master lies unconscious in the corner. The dog, taken in as he is by the thief, certainly thinks his master is in the room, and he is right (although, of course, he is wrong about this man's being his master). But just as clearly, the dog does not know that his master is in the room. We require conclusive reasons even in the case of animals. Would the dog have smelled what he did, and heard what he did, if his master was not in the room? If he *would have*, or *might have*, then he doesn't know.¹⁶

15 This seems to be the essence of Arthur Danto's argument (against Hintikka) that "S knows that P" does not entail "S knows that S knows that P" in "On Knowing that We Know," *Epistemology: New Essays in the Theory of Knowledge*, Avrum Stroll (ed.), 1967, pp. 49–51.

16 Robert Sleigh has suggested an interesting modification of this case – one that would appear to cause difficulty. Suppose that the dog *is* taken in by the thief but, in addition, circumstances are such that the thief *would not* be in the room *unless* the dog's master was also there. It may be a bit difficult to imagine circumstances of this sort, but other examples (not involving animals as the knowers) can easily be manufactured (Chisholm's example of a man mistaking a sheep dog for a sheep in a field works nicely when the further condition is added that the dog would not be there unless there were some sheep in the field). With such a modification it would seem that the dog has conclusive reasons for believing his master present since it would not have heard and smelled what it did unless his master was present. But does the dog know that his master is present? I think it natural to read this situation in such a way that the dog believes that his master is in the room *because* he mistakenly believes that this man (the thief) is his master. If read in this way, of course, there is no difficulty since the dog's basis for believing that his master is present is obviously defective – it is not true, nor does the dog know, that this man is his master. If, however, we read this case in such a way that the dog simply takes the sounds and smells as a sign of his master's presence (without misidentifying anyone), then the dog *does* have conclusive reasons (he would not have smelled and heard what he did unless his master was present), but I should want to say that in this case *he knows* that his master is present. I do not think that this is an excessively strained interpretation. It seems to me quite similar to situations in which we know that something is so on the basis of some indicator or sign (e.g., an instrument reading) but are ignorant as to the mechanism *in*