

Dretske and His Critics

Edited by
Brian P. McLaughlin

Basil Blackwell

Copyright © Basil Blackwell 1991
 Copyright © "Ways of Establishing Harmony", Daniel C. Dennett 1991
 First published 1991

Basil Blackwell, Inc.
 3 Cambridge Center
 Cambridge, Massachusetts 02142, USA

Basil Blackwell Ltd
 108 Cowley Road, Oxford, OX4 1JF, UK

All rights reserved. Except for the quotation of short passages for the purposes of criticism and review, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

Except in the United States of America, this book is sold subject to the condition that it shall not, by way of trade or otherwise, be lent, re-sold, hired out, or otherwise circulated without the publisher's prior consent in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

Library of Congress Cataloging in Publication Data

Dretske and his critics/edited by Brian McLaughlin.

p. cm.

Includes bibliographical references and index.

ISBN 1557-86198-6

1. Dretske, Fred I. 2. Knowledge, Theory of—History—20th century. 3. Content (Psychology)—History—20th century.

4. Psychology—Philosophy—History—20th century. 5. Visual perception—History—20th century. I. McLaughlin, Brian P.

B945.D744D74 1991

121'.092—dc20

90-46083

CIP

British Library Cataloguing in Publication Data

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

WID-LC

B

945

.D744

D74

1991

Typeset in 10 on 11 pt Times
 by Graphcraft Typesetters Ltd., Hong Kong.
 Printed in Great Britain by Billing & Sons Ltd, Worcester

1111

Skepticism, Relevance, and Relativity

Stewart Cohen

from Brian McLaughlin, ed.

Dretske & His Critics (Blackwell, 1991)

The relevant alternatives response to skeptical arguments is largely due to the work of Fred Dretske.¹ The primary virtue of this approach is that it can accommodate two opposing strands in our thinking about the concept of knowledge. As Peter Unger has observed, knowledge is an *absolute* concept.² On one interpretation, this means that the justification or evidence one must have in order to know must be absolute or conclusive³ – one's evidence must eliminate all alternatives to what one believes (where an alternative is a proposition incompatible with what one believes). When there is an alternative to what one believes that one's evidence does not eliminate, there is a strong intuitive pull in favor of saying that one does not really know what one believes. This aspect of our thinking about knowledge is precisely what skeptical arguments exploit. Calling our attention to alternatives that our evidence can not eliminate (e.g., hallucinations, dreams) the skeptic shows that this requirement of conclusive reasons is seldom, if ever, satisfied.

This conclusion flies in the face of our deeply entrenched supposition that we do know the truth of many ordinary empirical propositions. Thus, there is a tension between two components in our thinking about knowledge. We believe that knowledge is, in the sense indicated, an absolute concept and yet we also believe that there are many instances of that concept.

There would seem to be two options, each involving the denial of one of the components on the basis of the other. Peter Unger argues from the absolute character of knowledge to a skeptical conclusion, apparently finding absoluteness to be too central a component of our concept of knowledge to be relinquished.⁴

Most philosophers have taken the other course, choosing to respond to the conflict by giving up, perhaps reluctantly, the absoluteness criterion. The so-called fallibilists argue that this is precisely what we need to do, holding as sacrosanct our common-sense belief that we know many things.⁵ Each approach is subject to the criticism that it preserves one aspect of our ordinary thinking about knowledge at the expense of denying another.

Dretske's theory of relevant alternatives offers an elegant and ingenious way out of this unsatisfactory situation. It provides a view of knowledge that respects its absolute character as well as our commonsense suppositions regarding our epistemic accomplishments. What Dretske proposes is that we qualify, rather than deny, the absolute character of knowledge. According to Dretske, knowledge is relationally absolute – "absolute, yes, but relative to a certain standard."⁶ In order to know a proposition to be true, our evidence need not eliminate all the alternatives to that proposition. Instead we can know when our evidence eliminates all the *relevant* alternatives (where the set of relevant alternatives, a subset of the set of all alternatives, is determined by some standard). As such, there is no skeptical mileage to be gained from the absolute character of knowledge. The fact that the skeptic can discover alternatives to the propositions we claim to know, that our evidence cannot eliminate, does not in itself guarantee a skeptical result. For knowledge, being relationally absolute, requires that our evidence eliminate only the relevant alternatives. And the standards for knowledge are such that skeptical alternatives fail to be relevant.

Dretske motivates this view of the concept of knowledge by noting that other concepts exhibit the same logical structure. Two of his favored examples are the concept of empty and the concept of flat. Both appear to be absolute concepts – a space is empty only if it does not contain anything and a surface is flat only if it does not have any bumps. However, as Dretske points out, the absolute character of these concepts is relative to a standard. In the case of flat, there is a standard for what counts as a bump and in the case of empty there is a standard for what counts as a thing. We would not deny that a table is flat because a microscope reveals irregularities in its surface. Nor would we deny that a warehouse is empty because it contains particles of dust. To be flat is to be free of any relevant bumps. To be empty is to be devoid of all relevant things.

Dretske suggests that we use the logical structure of these concepts as a model for understanding the concept of knowledge. Thus Dretske proposes that to know a proposition is (*inter alia*) to have evidence that eliminates all relevant alternatives.

I believe that Dretske's view of knowledge is, in most respects, correct.⁷ In particular I am sympathetic to Dretske's claim that the relational character of knowledge is a kind of contextual relativity or context sensitivity. I propose to show that if we understand this claim in the right way, the theory of relevant alternatives provides a powerful response to the skeptical challenge.

The Skeptical Paradox

Initially one might think that the theory of relevant alternatives begs the question against skepticism. After all, precisely what the skeptic claims is

that the existence of alternatives our evidence cannot eliminate undermines our claims to know. The relevant-alternatives theorist merely responds that the alternatives the skeptic has invoked are not relevant, i.e., do not have to be eliminated in order for us to know. Since this is the point at issue, doesn't the relevant-alternatives theorist need an argument in support of this crucial claim?

Although in one sense, the relevant-alternatives theorist does not have an argument, he is able to appeal to our strong intuition that in many cases we do know things. And it is not apparent that the skeptic has an argument that undermines those intuitions.

One might think that there is such an argument available to the skeptic. The argument emerges once we are precise about what Dretske means when he refers to our evidence *eliminating* an alternative. According to Dretske, one's evidence eliminates an alternative just in case one's evidence against the alternative is good enough for one to know the alternative is false.⁸ But then the skeptic can claim, quite plausibly, that whatever else we say about the significance of skeptical alternatives, it is not correct to say that we know they are false. We might think that we have some reason to believe that we are not deceived in the ways the skeptics suggest we might be, but it is very hard to hold that we *know* we are not so deceived. One of Dretske's own examples of this is the counter-intuitiveness of claiming that when we are viewing zebras at the zoo, we *know* that we are not in fact seeing cleverly-disguised mules.⁹ While we do have some evidence against the possibility of such a deception, intuitively it is not strong enough for us to *know* that we are not so deceived.

At this point the skeptic can appeal to an epistemic closure principle with considerable intuitive force:

If *s* knows *q*, and *s* knows that *q* entails not-*h*, then *s* knows not-*h*.¹⁰

It follows from this principle and the just-cited claim that we fail to know the falsity of skeptical alternatives, that we fail to know the propositions we ordinarily claim to know. (Since those propositions trivially entail the falsity of skeptical alternatives, e.g., "one sees a zebra" entails that "one does not see a cleverly-disguised mule.")

This argument presupposes that every alternative is relevant. For the closure principle is just a precise expression of the unqualified claim that knowledge is absolute (requires the elimination of every alternative). But Dretske accepts only the claim that knowledge is relationally absolute. Indeed, Dretske has appealed to the fact that we know ordinary empirical propositions and the fact that we do not know the falsity of skeptical alternatives as an argument against the truth of the closure principle.

Moreover, some have agreed with the skeptic, contrary to Dretske, that the closure principle is true. But, against the skeptic, they use the closure principle along with the claim that we do have knowledge, to

reject the claim that we do not know the falsity of skeptical alternatives.¹¹ So is the skeptic begging the question against these other positions?

It is not clear how to assess this situation. I would suggest that what we are confronting here is a paradox – a set of inconsistent propositions each of which has considerable independent plausibility:

- 1 We know that some ordinary empirical propositions are true.
- 2 We do not know that skeptical alternatives are false.
- 3 If *s* knows *q*, and *s* knows that *q* entails not-*h*, then *s* knows not-*h*.

Because each member of the set has independent plausibility it would seem arbitrary and unsatisfying to appeal to any two members of this inconsistent triad as an argument against the third. Such a strategy does not provide what any successful resolution of a paradox should, viz., an explanation of how the paradox arises in the first place. As such, it is a constraint on any resolution that defends common sense against the skeptic, that it explain the appeal of skeptical arguments. For it is that very appeal that gives rise to the paradox. While initially we think we can claim to know unproblematically, under pressure from skeptical arguments, we begin to worry whether we really do have knowledge. Often we find ourselves vacillating between thinking we have knowledge and worrying that we do not. As Dretske notes, skeptical arguments have a “persisting and undiminished appeal.”¹² No successful response to the skeptical paradox can fail to explain that appeal.

Contextualism

Can the theory of relevant alternatives explain the appeal of skeptical arguments? The core idea of the relevant alternatives approach is that knowledge despite its absolute character is still in some sense a relative notion. Dretske tells us that “this absolute notion [knowledge] exhibits a degree of contextual relativity in its ordinary use.”¹³ What does it mean to say that the relativity of knowledge is *contextual*? For Dretske the idea seems to be this. One knows relative to a set of relevant alternatives – knowledge requires the elimination of all and only relevant alternatives. And the set of relevant alternatives is determined relative to a context. In this way, knowledge is a context-relative or context-sensitive notion.

There is more than one phenomenon that could be referred to by the expression “contextual relativity” and thus more than one phenomenon that could be described by saying that the set of relevant alternatives is determined relative to a context. Whether or not the theory of relevant alternatives can discharge its burden of explaining why we are troubled by skeptical arguments will depend on how we construe Dretske’s crucial notion of contextual relativity.

We can begin by examining Dretske’s own account of how the theory

of relevant alternatives accomplishes this task. He tells us that with the theory,

we get a better perspective from which to understand the *persisting* and *undiminished* appeal of skeptical arguments. Most philosophers have experienced the futility of trying to convince a devoted skeptic, or just a newly converted freshman, that we do know there are tables and chairs despite the possibility of dreams, hallucinations, cunning demons and diabolical scientists who might be toying with our brain on Alpha Centuri (Nozick’s example). Somehow, in the end, we seem reduced to shrugging our shoulders and saying that there are certain possibilities that are just too remote to worry about. Our evidence isn’t good enough to eliminate these wider hypotheses because, of course, these wild hypotheses are carefully manufactured so as to neutralize our evidence. But dismissing such hypotheses as too remote to worry about, as too fanciful to have any impact on our *ordinary use* of the verb “to know,” is merely another way of saying that for purposes of assessing someone’s knowledge that this is a table, certain alternative possibilities are simply not relevant. We are doing the same thing (or so I submit) as one who dismisses chalk dust as irrelevant, or too insignificant, to worry about in describing a classroom as empty. What it is important to realize, especially in arguments with the skeptic, is that the impatient dismissal of his fanciful hypotheses is not (as he will be quick to suggest) a mere practical intolerance, and refusal to confront, decisive objections to our ordinary way of talking. It is, rather, a half-conscious attempt to exhibit the *relationally* absolute character of our cognitive concepts.¹⁴

This passage seems not to deliver what it promises. If “our cognitive concepts” are such that “certain possibilities are just too remote to worry about,” if these possibilities are “too fanciful to have any impact on our ordinary use of the verb ‘to know’” then wherein lies “the persisting and undiminished appeal of skeptical arguments?” Why don’t we immediately respond to skeptical arguments by objecting that skeptical alternatives are too remote and fanciful to undermine our knowledge? It is true that we sometimes are inclined to do just that. But the skeptical problem arises precisely because we can not always sustain that attitude. Sometimes we begin to worry that our inability to eliminate these alternatives does threaten our claims to know. Again, often we vacillate. The appeal of skeptical arguments is indeed persisting and often remains undiminished despite our best attempts to rebuff them. But I see nothing in this passage that shows how the theory of relevant alternatives enables us to understand why this is so.

Nonetheless, I believe the theory does have the means to provide such an explanation. It is here that we must exploit the notion of contextual relativity. What is required is that we distinguish between two very

different phenomena that could be thought of as the contextual relativity of knowledge.

One characteristic of knowledge that we might think of as contextual relativity is the fact that whether a subject has knowledge depends on certain features of the subject's circumstances beyond the evidence the subject has. This phenomenon has been illustrated by numerous cases in the literature where various factors such as the availability of unpossessed misleading counter-evidence, the evidential beliefs of the members of relevant social groups, and the objective probability of being wrong, undermine the knowledge of a subject who otherwise possesses adequate evidence.¹⁵ Thus, two subjects could possess the same evidence for the truth of a certain proposition, and one of them know the proposition while the other fails to know the proposition. In this sense, whether a subject knows on the basis of certain evidence is determined relative to the context, i.e., relative to the extra-evidential circumstances of the subject. Dretske refers to cases of this kind and claims as a virtue of the relevant-alternatives theory that it can explain them.¹⁶ In essence, he claims that these factors are part of what determines the set of relevant alternatives. For example, in normal circumstances, when one sees a Gadwall duck, the alternative that instead one sees a look-alike Grebe is not relevant – one's evidence does not have to eliminate it in order for one to know that one sees a Gadwall. However, if the circumstances happen to be such that there are look-alike Grebes in the immediate vicinity, then the alternative that one sees such a Grebe is relevant. Thus one subject may know that he sees a Gadwall on the basis of visual evidence while another fails to know on the basis of the same evidence, due to the fact that the extra-evidential circumstances play a role in determining the relevance of alternatives.

This kind of contextual relativity, while in one sense interesting, is in another sense trivial. That knowledge is relative to the extra-evidential circumstances of the subject follows, on most views, from the truth condition for knowledge alone. On any fallibilist view, two subjects in different cases could have equivalent evidence for a proposition and yet only one of them know the proposition merely because in only one case is the proposition true. So to claim that knowledge is context-relative in this sense, is not in itself to make a novel claim. Of course one of the benefits of the extended Gettier controversy is that it has produced a variety of novel types of extra-evidential circumstances besides the truth of the proposition, that play a role in determining whether a subject knows the proposition. What is novel is not that knowledge is, in this sense, context-relative. Rather, what is novel is the myriad of circumstances that have been discovered to be part of this extra-evidential context.

Another characteristic of knowledge that could be referred to as contextual relativity or context sensitivity is more accurately described as a characteristic of attributions of knowledge, viz., indexicality. To say that attributions of knowledge are indexical is to say that the truth-value of an attribution of knowledge is determined relative to the context of *attribu-*

tion, i.e., relative to the speaker or the conversational context.¹⁷ Applied to the theory of relevant alternatives, the view would be that the purposes, intentions, presuppositions, etc., of attributors of knowledge – speakers and listeners – play a role in setting the standards of relevance.¹⁸ This kind of relativity is quite different from relativity to the extra-evidential circumstances of the subject of the attribution.¹⁹ If attributions of knowledge are context-relative in this sense, then two speakers could simultaneously say of a subject, "S knows p" and it be the case that what one speaker says is true while what the other says is false. Moreover, holding the circumstances of S fixed, one speaker may say "S knows p" while another says "S does not know p" without it being the case that they have contradicted each other. This is because the sentence "S knows p" will have different truth conditions in different contexts of attribution. If "know" is an indexical, then it will express different knowledge relations in different contexts of attribution and thus sentences containing "know" will express different propositions. Applied to the relevant alternatives view, in different contexts of attribution, attributions of knowledge will involve different standards of relevance.

For example, one factor that Dretske suggests will affect the relevance of an alternative is its remoteness.²⁰ Suppose then that A attributes knowledge to S that S sees a Gadwall.²¹ The context of attribution will determine a standard of remoteness, viz., a standard for when an alternative is too remote to be relevant. That attribution involving that standard is then evaluated at the circumstances in which S sees the Gadwall, where here the pertinent factor is the actual remoteness in those circumstances of the alternative that S sees a look-alike Grebe. And this will depend on whether there are look-alike Grebes in the area – their number, frequency, etc. Thus, the truth-value of an attribution of knowledge is determined in part, by whether the circumstances of the subject are such that there are any alternatives which are not too remote, relative to the standard of remoteness yielded by the context of attribution. The truth-value of a knowledge attribution will vary as we vary either the context of attribution that determines the standard of relevance or as we vary the extra-evidential circumstances governed by the standards.

As I read Dretske, he discusses both of these distinct phenomena under the rubric of contextual relativity. However, it is crucial to distinguish between them. It is relativity to the context of attribution (indexicality) that is suggested by Dretske's analogy with words like "flat" and "empty." Moreover, it is this kind of contextual relativity (indexicality) that the relevant alternatives theory requires if it is to provide an explanation of the appeal of skeptical arguments.

Consider Dretske's treatment of "flat":

For although nothing can be flat if it has *any* bumps and irregularities, what counts as a bump or irregularity depends on the type of surface being described . . . a road can be perfectly flat even though one can *feel* and *see* irregularities in its surface, irregularities which,

continue to govern their attributions (perhaps only among themselves) even if they were to return to Kansas.

A further example shows that it is the context of attribution that determines the standards that apply (rather than the location of the subject of the attribution). Beings that vary enormously in size would most likely differ in the standards of flatness they employ. A giant would differ from us in his attributions of flatness. Surely there is no basis for us to say that his attributions of flatness are incorrect. If we were to claim that his size makes him insensitive to irregularities that really exist, then our own judgments would be subject to the same objection from ant-sized beings. Again the context of attribution will be what selects the standard. This holds true regardless of where the particular road is located.

Dretske's remarks about "empty" suggest that we should view this family of terms as indexical in this way:

Something is empty (another absolute concept according to Unger) if it has nothing in it, but this does not mean that an abandoned warehouse is not really empty because it has light bulbs or molecules in it. Light bulbs and molecules do not count as *things* when determining the emptiness of warehouses. For purposes of determining the emptiness of a warehouse, molecules (dust, light bulbs, etc.) are irrelevant. This isn't to say that, if we changed the way we used warehouses (e.g., if we started using, or trying to use, warehouses as giant vacuum chambers), they *still* wouldn't count. It is only to say that, given the way they are now used, air molecules (dust particles, etc.) don't count.²⁴

So whether a space is empty depends on the purpose for which we use it. That is to say our criteria for determining whether a space is devoid of things, in the relevant sense, depend on our purposes, goals, intentions, etc. And as Dretske notes, these facts about us may vary. We might decide to use a space that we had previously used as a warehouse, as a vacuum chamber. By so deciding, we change the criteria for determining what counts as a thing in the relevant sense and we thus change our criteria for the correctness of attributions of emptiness. As Dretske says in another work, "The concept [emptiness], though absolute, has a built-in plasticity (in the idea of a 'relevant' thing) that is responsive to the interests and purposes of people applying it."²⁵

So we could imagine two speakers, one of whom says "The room is empty," since he intends to use it for storage, whereas another speaker says "The room is not empty" since he intends to use it as a vacuum chamber. Ostensibly, the two speakers make conflicting attributions. But since the contexts of attribution involve different standards, their attributions do not really conflict. Both statements could be true.

If this account of the semantics of "flat" and "empty" (or of the concepts of flat and empty) is correct, then one can not determine the

were they to be found on the surface, say, of a mirror would mean that the mirror's surface was not really flat.²²

Although, *in general*, I agree with Dretske's remarks, I think the situation is more complicated than he implies. What counts as a bump (the standard of flatness) can vary even with respect to one type of surface. Suppose it is given that what is being described is a road (rather than a mirror). Smith who is from Colorado says "The road is flat." Jones, from Kansas, says "The road is not flat." Their disagreement stems from the different standards they employ in making attributions of flatness. Jones, being from Kansas has much stricter standards than someone from Colorado, like Smith. Now suppose we ask who is right and who is wrong about the road. Surely it would be inappropriate to decide the issue in favor of one against the other. They make their attributions relative to different standards. Relative to its respective standards, each attribution could be correct (or incorrect). There is no contradiction because by their use of different standards they have, in effect, said different things about the road.

The important thing to notice is that there is nothing in the circumstances on the basis of which we could say that either Smith's or Jones's standards are, in some absolute sense, the *correct* standards – the standards that determine the truth-value of the attribution. I do not mean to imply that it is the intentions of the speaker alone that determine which standards are in effect. Factors such as the expectations of the listeners and the presuppositions of the conversation can play a role in the determining the standards.²³ The point is that the standards in effect result from the context of attribution (the intentions, presuppositions, etc., of speakers and listeners) rather than from the circumstances of the subject (in this case a road) in which the attribution is evaluated.

It might seem as if the standards are determined by such circumstances. In particular, one might claim that it is the location of the road that determines which standards are in effect. So in Colorado, looser standards apply than in Kansas. However, if a group of Kansas residents are traveling through Colorado, they might all agree that they have yet to encounter a flat road. Surely we should not insist that they have all agreed on something false. As residents of Kansas, they assess the flatness of a road relative to stricter standards. And why should we evaluate what they say relative to standards other than the ones they as a group share and understand one another to be using? What they say does not contradict what a group of Colorado residents might agree upon if they were to utter the sentence "The road the Kansas residents are on is flat." Relative to the standards that operate in their respective contexts of attribution, both claims can be true (or false).

If the Kansas residents spend enough time in Colorado, they may be led to relax their standards. Were that to happen, they may say truly of a road in Colorado, "That is a flat road." These same standards could

truth-value of sentences containing these terms without considering the context of attribution (or utterance), viz., the intentions, goals, interests, etc., of the attributors. That is we must view the use of these terms as involving an indexical reference to standards or criteria.

If our treatment of "flat" and "empty" is to serve as a model for "know" (or if our treatment of the concepts of flat and empty is to serve as a model for the concept of knowledge) then we should view "know" as indexical in this way. Indeed Dretske at times seems to take this view explicitly:

When a possibility becomes a relevant possibility is an issue that is, in part at least, responsive to the interests, purposes, and yes, values of those with a stake in the communication process. The flow of information, just like the cognitive exploits it makes possible, is a process that exhibits some sensitivity to the variable purposes of those who send and receive this information.²⁶

So whether a subject knows depends on the set of relevant alternatives, and the composition of that set is a function of the context of attribution (or, we could say, the context of communication). That is to say, the criteria of relevance will vary as the interest, purposes, etc., of those involved in making knowledge attributions vary.

This observation provides the key to explaining the appeal of skeptical arguments. Recall that Dretske claims that "for the purposes of assessing someone's knowledge that this is a table, [skeptical alternatives] are simply not relevant." I argued that this approach cannot, in itself, explain what Dretske acknowledges to be, the persisting appeal of skeptical arguments. To explain this appeal, we need to exploit Dretske's claims about the contextual relativity of knowledge. Dretske appeals to contextual relativity in an attempt to defuse skeptical arguments, but it appears that he does not clearly distinguish between what I have claimed are two different kinds of contextual relativity:

One of the ways to prevent this slide into skepticism is to acknowledge that although knowledge requires the evidential elimination of all relevant alternatives (to what is known), there is a shifting, variable set of relevant alternatives. It may be that our birdwatcher does know the bird is a Gadwall under normal conditions (because look-alike grebes are not a relevant alternative), but does not know this if there is a suspicion, however ill-founded it may be, that there exist look-alike grebes within migrating range.²⁷

Here, Dretske cites an example of what I have referred to as relativity to the extra-evidential circumstances of the subject – in addition to the strength of the evidence the subject happens to possess, the context (extra-evidential circumstances) in which the subject possesses that evidence can affect whether the subject knows. If the issue concerns whether

someone knows that he sees a Gadwall duck, conditions like the actual presence of look-alike Grebes or the mere suspicion of their presence by a nearby ornithologist can affect whether a subject knows, by affecting the relevance of the alternative that he sees a look-alike Grebe. This makes it look as if Dretske's response to the skeptic is to point out that while in rare circumstances, skeptical alternatives such as look-alike Grebes and cleverly disguised mules may be relevant, normally this is not the case. Presumably Dretske would think that circumstances are never such that the more radical skeptical alternatives (e.g., brain-in-a-vat, or Cartesian demon scenarios) are relevant. As he says earlier in the article, "for the purposes of assessing someone's knowledge . . . certain alternative possibilities are simply not relevant." Again this response does not explain the persisting appeal of skeptical arguments.

But Dretske goes on to say with respect to this effect of the circumstances of the subject on the relevance of alternatives for knowledge,

This will (or should) be no more unusual than acknowledging the fact that a refrigerator could truly be described as empty to a person looking for something to eat, but *not* truly described as empty to a person looking for spare refrigerator parts.²⁸

Here, however, there is a crucial shift in the phenomenon at issue. The examples concerning attributions of emptiness do not show that the correctness of such attributions depends on the circumstances of the subject of the attribution – in this case a refrigerator. Rather, they show that the correctness of an attribution of emptiness depends on the context of attribution, viz., the intentions, purposes, etc., of the attributor. This is contextual relativity in the sense of indexicality. In this sense, two speakers could simultaneously say of the same refrigerator "It is empty," and it be the case that what one says is true and what the other says is false. This could occur if, as in Dretske's example, one speaker is looking for something to eat, while another is looking for spare refrigerator parts. The two different contexts yield different criteria for when the presence of a certain type of thing is relevant to attributions of emptiness.

It is this latter kind of contextual relativity that allows us to explain the appeal of skeptical arguments. We have noted that simply to say that skeptical alternatives are, in some sense, too remote to be relevant is to fail to do justice to the apparent threat that they present to our knowledge claims. While we often feel that skeptical alternatives are too remote to threaten our knowledge claims, at other times we find them quite troubling. When we think about skepticism, we often vacillate with respect to the relevance of these alternatives, and thus with respect to whether we know. By supposing that the criteria or standards of relevance, e.g., standards of remoteness or probability, are relative to the context of attribution (the intentions, interests, purposes, etc., of attributors), we can explain our tendency to vacillate in this way. (In the same way, by supposing that the standards for flatness are relative to the

context of attribution, we can explain why we might vacillate over whether to describe a particular road as flat).

In everyday contexts, our standards are such that, normally, skeptical alternatives are not relevant. In rare circumstances, some skeptical alternatives will be relevant according to everyday standards, e.g., if we are looking at a zebra in an exhibit that actually contains mules cleverly disguised to look like zebras. Our confidence in the truth of our everyday knowledge attributions is explained by the fact that normally, skeptical alternatives will not be relevant according to the standards in effect. However, when we are confronted with skeptical arguments, we may come to consider skeptical alternatives as relevant, thereby shifting our standards. Skeptical arguments are forceful, precisely because they can have this effect on us.²⁹ In this new context, where the standards for knowledge are stricter, attributions of knowledge which would be true in everyday contexts, are false. But we are not thereby constrained to use skeptical standards for our knowledge attributions. Upon further consideration, we may decide that skeptical alternatives are too remote to count as relevant, thereby shifting the standards once more. Again, sometimes we vacillate between considering skeptical alternatives as relevant and dismissing them as irrelevant. By supposing that attributions of knowledge are sensitive to context in this way, we do justice both to our strong inclination in everyday life to say we know things and the "persisting and undiminished appeal of skeptical arguments."³⁰

One might object that the relevant alternatives view understood in the way I have been suggesting does not really constitute a reply to skepticism. For one could respond that what the skeptic denies is that we ever know relative to standards that make every alternative relevant. Thus to argue that we know only relative to standards that limit the set of relevant alternatives is not to address the skeptic at all.³¹

The problem with this objection is that it gets things backwards. If the theory of relevant alternatives is correct, it is the skeptic who does not really address our everyday knowledge attributions. After all, the interesting issue is whether the claims to know that we make in everyday life are true. What is initially startling about skeptical arguments is that they seem to call those claims into question. The theory of relevant alternatives shows how those claims – properly interpreted – are true. We do know relative to the standards that govern our everyday knowledge attributions.

This is not to deny that knowledge attributions involving stricter standards are false. Rather, the point is that those stricter standards do not ordinarily apply to our knowledge attributions.

The Closure Principle

Let us return briefly to the paradox that is the crux of the skeptical problem. Construing the relevant alternatives theory as I have suggested,

it remains unclear which member of the inconsistent triad we should deny. Dretske has argued that proposition (1), the closure principle, is false, since we can know a proposition is true without knowing the falsity of certain alternatives.³² For example, we can know that we see a zebra without knowing that it is not the case that we see a cleverly disguised mule.

However, if attributions, of knowledge are sensitive to the context of attribution, then it is not clear that the theory of relevant alternatives entails the falsity of the closure principle. It has been argued that we must interpret the closure principle as applying within a fixed context of attribution. Moreover, relative to a fixed context, the principle is true.³³ So in contexts where, e.g., we know we see a zebra, we also know that we do not see a cleverly disguised mule, and in contexts where we fail to know that we do not see a cleverly disguised mule, we do not know that we see a zebra.

On this view, we can argue that the skeptical paradox arises because of our failure to pay attention to contextual shifts. The closure principle holds relative to a context, and one of the other two propositions that constitute the paradox will be false, depending on the standards of the particular context. In no context are they both true.

Of course, I have only given here a sketch of such a view and I certainly have not shown that Dretske is wrong to interpret the relevant alternatives view as entailing the falsity of the closure principle. However, a discussion of the details of this issue would take us too far afield.³⁴

A Further Application

I have been arguing that if we construe attributions of knowledge as sensitive (or relative) to the context of attribution, we have available an effective strategy for dealing with skeptical arguments. But the utility of this view goes beyond its application to skeptical puzzles.³⁵ A further illustration is provided by a puzzle that Dretske discusses. It concerns the role of the importance of what is known in determining the set of relevant alternatives. It is crucial to note that what is at issue is, as Dretske says, the role of the importance of what is known for *speakers and listeners* in determining what alternatives are relevant. Dretske discusses an alleged example of this phenomenon:

The fuel gauge (and associated mechanism) that suffices for knowing that you still have some gasoline (when driving in the city) is just not good enough for knowing that there is sufficient liquid coolant surrounding the reactor [on Three Mile Island]. This somewhat paradoxical fact (the fact, namely, that a particular instrument should be good enough to give knowledge in one place, not good enough in another) is to be explained, some would say, by the fact that as the stakes go up, the stakes associated with being right about

what one purports to know, so does the size of the relevancy set. There are more possibilities that must be eliminated in the nuclear power plant than must be eliminated in the automobile. In particular, a malfunction in the instrument itself must be guarded against in the dangerous situation. If it isn't, one doesn't know.³⁶

Although in this passage, Dretske notes that it is paradoxical "that a particular instrument should be good enough to give knowledge in one place, not good enough in another," he concedes that "there is some appeal to this point. . . ." However, he claims that the point is mistaken.

I see no reason why a standard automobile gauge, transplanted from the automobile to the nuclear power plant, functioning as the only indicator of coolant level, should not, assuming it continues to function reliably (as reliably as it did in the automobile), be able to do precisely what the more expensive instruments do - viz., tell the operators that the coolant is at a safe level.³⁸

Since Dretske admits there is some appeal to the point, i.e., some intuitive pull toward treating the two cases in the example differently, if he claims that the point is mistaken, he must provide some explanation for our intuitive response. What Dretske says is that our intuitions are indeed responding to a difference between the two cases. There is an uncertainty that is appropriate in the reactor case that is not appropriate in the automobile case. It is just that the uncertainty concerns not whether the reactor operator knows what the coolant level is, but rather "when (due to gauge malfunction) they stop knowing it."³⁹

I find Dretske's treatment of this case to be puzzling. Presumably, his reservation about the original point is that, as he says, it is paradoxical to claim that a gauge good enough for knowing in one place is not good enough in another. But then why should it not be equally paradoxical that the same change of location can make a difference in whether we should be uncertain about when a subject who knows, stops knowing. If the source of our intuitive response is as Dretske says, then he must believe (as I do) that we do not have the same uncertainty about when we will stop knowing that we have enough gas on the basis of the gauge in the automobile.

Moreover, the point cannot just concern our uncertainty as to when, due to gauge malfunction, we stop knowing that there is enough coolant. If the reliability of the gauge over time is at issue, then there will come a time at which we should be uncertain over whether we have stopped knowing, i.e., whether we still know. (We can assume for the purposes of this example that there is a significant lag time between when the fuel level drops and when it is detectable by means other than the gauge).⁴⁰

I agree with Dretske that our intuitive responses to examples of this

kind can be puzzling. However, we can resolve the puzzle by viewing the knowledge attributions we make as sensitive to the context of attribution. Recall that what is paradoxical about the example is that when we inquire whether a subject knows on the basis of a gauge, our intuitions pull us toward different answers depending on where the gauge is located (an automobile as opposed to a nuclear reactor). But it is paradoxical only on the assumption that we are asking the same question in each case. And if attributions of knowledge involve an indexical reference to a standard, this assumption may be incorrect. In different contexts of attribution, different standards can apply, depending on the intentions, values, etc., of the attributor. This is crucial for our present purposes since the original issue that Dretske raises is whether the importance for *speakers and listeners* of what is known can affect the composition of the relevancy set. If we suppose that the importance of what is known for attributors (speakers and listeners) does affect the standards of relevance that operate in a context, we get an explanation for what is otherwise puzzling. It is not paradoxical that our intuitions may lead us to judge that a subject can know on the basis of a particular instrument in one place even though our intuitions lead us to hesitate in making the same judgment when the instrument is in another place, if our viewing the stakes as higher in the latter case leads us to employ stricter standards. Moreover, this is just what we should expect if as Dretske says, the composition of the relevancy set is "in part at least responsive to the interests, purposes, and . . . values of those with a stake in the communication process."⁴¹

A further example shows that attributors may employ different standards of relevance for the same case, if they diverge in what they consider to be important: Smith is very worried about whether the cruise ship which he and Jones are on will sink. Jones is not at all worried about this, having the utmost confidence in the seaworthiness of the ship. Smith asks Jones if she knows whether there are enough life preservers aboard. Jones replies that she knows there are a sufficient number since the tour guide says that the ship carries at least one life preserver per passenger. Smith, given his concern about the ship going down, may hesitate to agree with Jones that she knows. Because of the increased importance he attaches to the presence of the life preservers, Smith is only willing to attribute knowledge to Jones (or himself) if Jones has actually seen the life preservers on the ship. That is to say, by Smith's standards the possibility of misinformation in the tour guide is a relevant alternative. Jones, because she does not attach much importance to the presence of the life preservers is willing to attribute knowledge to herself on the basis of what the tour guide says. According to her standards, the alternative that the tour guide is misleading is not relevant. And, of course, there is no issue concerning what the correct standards are absolutely, if we take context-sensitivity seriously. Like attributions of flatness and emptiness, the standards that apply for attributions of knowledge will vary with the context.

Criteria of Relevance

What makes an alternative relevant? While philosophers have been able, through the use of examples, to point to some considerations that play a role in determining relevance, it has proven to be very difficult to give precise formulations of general criteria of relevance.

For example, what is frequently discussed as a major consideration is the probability or remoteness of an alternative. Dretske illustrates this phenomenon in his case of the Gadwall ducks and the look-alike Grebes. According to Dretske, if there are no such Grebes or if they exist but because of geographical barriers are confined to a distant area, then the alternative that someone sees such a Grebe is not relevant to whether they know they see a Gadwall. If they do exist and can migrate to the area, then this alternative is relevant. Dretske suggests that the "difference between a relevant and an irrelevant alternative resides . . . in the kind of possibilities that actually exist in the objective situation."⁴³ I am sure Dretske would agree that this is not a very precise formulation. Because there are so many complex and controversial variations in examples like this, it is exceedingly difficult to capture the distinction between relevant and irrelevant alternatives in a precise criterion.⁴⁴

Because of this problem, many philosophers have been led to have grave doubts about the adequacy of the theory of relevant alternatives as a response to skepticism.⁴⁵ Dretske, in his own discussion, holds that the anti-skeptical power of the theory of relevant alternatives (as well as other virtues of the theory) "can only be harvested if certain questions [including, "What makes an alternative relevant?"] can be given reasonable answers."⁴⁶ Apparently, Dretske either is more sanguine about the prospects for formulating precise criteria of relevance or he believes that something short of precise criteria constitutes a reasonable answer to the question.

I think that the latter option is the correct position to take. While I am not at all sanguine about the prospects for developing precise criteria, I think that the demand for such criteria in order for the theory to provide an adequate response to skepticism, is misguided. We can see this by looking at two ways we might construe the relevant alternatives approach.

One way for the relevant alternative theorist to argue would be to begin by appealing to criteria of relevance in the abstract. The claim would have to be that one can see simply by reflecting on the concept of relevance that the proposed criteria capture what it is about alternatives that make them relevant. Having gained assent on the correctness of the proposed criteria, the relevant-alternatives theorist would then proceed by arguing that when we apply these criteria to skeptical alternatives, we get the result that skeptical alternatives are not relevant. If we were to proceed in this way, then of course, the failure to provide a precise statement of the criteria would undermine the whole approach.

But we should not construe the relevant alternatives theorist as employing this argumentative strategy. Even if we did not worry about the precision problem, surely this strategy would be hopeless. We cannot pull criteria of relevance out of thin air. We formulate such criteria by reflecting on the deliverances of our intuition regarding cases. We try to devise criteria that capture those intuitions. We all have intuitions about relevance (the borderline cases notwithstanding). Since "relevance" is defined in terms of knowledge, our intuitions about relevance are at root, intuitions about knowledge.⁴⁷ Even the critics of the relevant alternatives approach demonstrate that they have such intuitions when they criticize various proposed criteria for yielding counterintuitive results.

So we should not construe the relevant alternatives theorist as arguing that skeptical alternatives are not relevant because the criteria show it. The irrelevance of skeptical alternatives is a datum provided by our intuitive judgments concerning what we know. It is illuminating to advert to various factors that account for relevance and it would certainly be desirable to have a precise formulation of the criteria of relevance. But surely it does not follow from our inability to provide a precise formulation, that we cannot legitimately apply the concept. Our inability to provide a precise formulation of the criteria of relevance derives from our inability to provide a precise formulation of the criteria of knowledge. If our inability to precisely formulate criteria made it illegitimate to apply a concept, the skeptic could establish his position simply by appealing to the fact that philosophers have been unable to formulate precise criteria of knowledge.

Of course if we want the theory of relevant alternatives to provide an analysis of knowledge in this sense, then the failure to provide precise criteria would constitute a failure of the theory. But this project is quite different from the project of responding to skepticism.

Again one might think that it begs the question against the skeptic to appeal to anti-skeptical intuitions about relevance. And there is a sense in which it does. But equally, the skeptic begs the question against common sense by appealing to his skeptical intuitions about relevance. In this sense, it is impossible for either side of the dispute not to beg the question against the other side.

We need to be very clear about the nature of the enterprise in which we are engaged. What we are confronted with is not an argument that forces us to be skeptics. We noted that just as the skeptic can combine the claim that we do not know that skeptical alternatives are false with the closure principle to conclude that we fail to know ordinary empirical propositions, one could combine the same closure principle with the claim that we do know ordinary empirical propositions to conclude that we know that skeptical alternatives are false. Or we could take Dretske's route and argue that the closure principle is false.

What we face is a paradox. We are inclined to accept each member of a set of propositions we know to be inconsistent. What we seek is a theory that rescues common sense from the skeptical worries presented by the

paradox. It is not a constraint on such a theory that it appeal to the skeptic. The project is not to demonstrate to the skeptic that we know. Rather it is to demonstrate to ourselves that we can claim to know without paradox.

One way to see that the problem of formulating criteria of relevance does not undermine the theory of relevant alternatives as a response to skepticism is to notice that the same problem ultimately confronts the skeptic. For just as the defender of common sense must explain the appeal of skepticism, the skeptic must explain the fact that we are strongly inclined to say that we know many things. The persistence of these intuitions (even after the skeptic has raised his alternative possibilities) is as problematic for the skeptic as the persisting appeal of skeptical arguments is for the anti-skeptic.

How can the skeptic explain away these intuitions? The most promising strategy is to claim that in cases where we think we know, while it is false that we know, nonetheless it is in some way appropriate or useful for us to say we know.⁴⁸ Thus, if there are no look-alike Grebes, although we still fail to know we see a Gadwall, there can be many purposes for which it is useful to distinguish such a case, from cases where there are lots of lookalike Grebes around. So according to the skeptic, our everyday pattern of knowledge attributions marks a distinction between cases where it is appropriate to say we know and cases where it is not. But in all cases, the skeptic maintains, it is literally false that we know.

If the skeptic argues in this way, and it is hard to see how he can avoid it, then he needs criteria of relevance as much as the anti-skeptic does. The only difference is that for the anti-skeptic, the criteria determine when an alternative is relevant to the *truth* of knowledge attributions, whereas for the skeptic, they determine when an alternative is relevant to the *appropriateness* of knowledge attributions. Since the skeptic needs criteria of relevance as much as the anti-skeptic does, our inability to be precise about the criteria should not incline us toward skepticism. If both sides of a dispute share a problem, the existence of the problem cannot favor one side of the dispute against the other.

Notes

Portions of this paper are taken from my paper "How to be a Fallibilist," in *Philosophical Perspectives* vol. 2, ed. James Tomberlin (Ridgeview Publishing, 1988). I would like to thank Jonathan Vogel and Scott Soames for valuable discussion.

1 See Fred Dretske's "Conclusive Reasons," *Australasian Journal of Philosophy*, 49 (May, 1971), pp. 1-22; "Epistemic Operators," *Journal of Philosophy*, 67 (December, 1970), pp. 1007-23; "The Pragmatic Dimension of Knowledge," *Philosophical Studies*, 40 (October, 1981), pp. 363-78; *Knowledge and the Flow of Information* (Cambridge, Mass.: MIT Press/Bradford Books, 1981); I will focus on "The Pragmatic Dimension of Knowledge." Also see Cohen, "How to be a Fallibilist;" Alvin Goldman, "Discrimination

and Perceptual Knowledge," *Journal of Philosophy*, 78 (1976), pp. 771-91; G. C. Stine, "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies*, 29 (1976), pp. 240-560.

2 See Peter Unger, *Ignorance: The Case For Skepticism* (Oxford: Oxford University Press, 1975); Peter Unger, *Philosophical Relativity* (Minneapolis: University of Minnesota Press, 1974).

3 See Dretske, "Conclusive Reasons."

4 See Unger, *Ignorance*; in *Philosophical Relativity*, Unger modifies his view. See Roderick Chisholm, *Theory of Knowledge* (Englewood Cliffs, NJ: Prentice-Hall, 1977); Gilbert Harman, *Thought* (Princeton, 1974); Keith Lehrer, *Knowledge and Justification* (Princeton, 1974).

6 "Pragmatic Dimension of Knowledge," p. 367.

7 I defend a view of this kind in "How to be a Fallibilist." In "Knowledge and Context," *Journal of Philosophy* (October 1986) and "Knowledge, Context, and Social Standards," *Synthese*, 73 (1987), I argue that this kind of view can be motivated independently of issues concerning skepticism.

8 "Pragmatic Dimension of Knowledge," p. 371.

9 "Epistemic Operators," p. 1016.

10 This closure principle may be subject to certain quibbles. But surely something very close to it is very intuitive. See Robert Nozick, *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press, 1981).

11 See Pollock, *Knowledge and Justification*; Stine, "Skepticism, Relevant Alternatives, and Deductive Closure;" Peter Klein, *Certainty: A Refutation of Skepticism* (Minneapolis: University of Minnesota Press, 1981).

12 "Pragmatic Dimension of Knowledge," p. 368.

13 *Ibid.*, p. 365.

14 *Ibid.*, p. 368.

15 See Goldman, "Discrimination and Perceptual Knowledge;" Harman, *Thought*; Lehrer, *Knowledge*.

16 Dretske says that he doubts whether some of these cases really are such that the subject fails to know. I discuss a case of this kind later in the text.

17 See David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic*, 8 (1979) pp. 513-43. Cohen, "Knowledge and Context;" Cohen, "Knowledge, Context, and Social Standards;" Cohen, "How to be a Fallibilist;" Dretske, "Epistemic Operators;" Dretske, "The Pragmatic Dimension of Knowledge;" Dretske, *Knowledge and the Flow of Information*; Goldman, "Discrimination and Perceptual Knowledge;" Stine, "Skepticism, Relevant Alternatives, and Deductive Closure."

18 Exactly how these features of the context determine the standards is a complex and difficult matter. But this is no special problem for the claim that attributions of *knowledge* are context-sensitive. The mechanisms of context-sensitivity (indexicality) are not very well understood in general. Even (relatively) uncontroversial cases of predicates whose application depend on context-sensitive standards face the same difficulty, e.g., flat and empty.

19 In cases of self-attributions of knowledge, we must distinguish between the same person as attributor and as subject of the attribution.

20 "Pragmatic Dimension of Knowledge," p. 371. Of course, "remoteness" is vague. Later in the text, I discuss the problem of our inability to be precise about the criteria of relevance.

21 Strictly speaking, since on the view we are considering, "knowledge" and "know" are indexicals, the discussion should be framed metalinguistically. So

- 108 I should say that A says "S knows that he sees a Gadwall" or A says "S has knowledge that he sees a Gadwall." Because the metalinguistic formulations are somewhat cumbersome, I will not always use them. But the reader should not be misled by this.
- 109 "Pragmatic Dimension of Knowledge," p. 366.
- 110 For a theory of how this works, see Lewis, "Scorekeeping."
- 111 "Pragmatic Dimension of Knowledge," p. 366.
- 112 *Knowledge and the Flow of Information*, p. 133.
- 113 *Ibid.*, p. 133.
- 114 "Pragmatic Dimension of Knowledge," p. 370.
- 115 *Ibid.*, p. 370.
- 116 For a more detailed account of how this occurs, see Cohen, "How to be a Fallibilist."
- 117 My point here is not that Dretske is unaware of this view. As we have seen, there are several places in his writings where he appears to endorse this view. (In addition see Dretske, "Epistemic Operators," p. 1022) I don't know whether Dretske would agree that our standards can shift so as to make skeptical alternatives relevant. He does not argue this way in his explanation of the appeal of skeptical arguments. He does say that the set of relevant alternatives will always be a proper subset of the set of alternative (in "Pragmatic Dimension of Knowledge.")
- 118 See Paule Yourgrau, "Knowledge and Relevant Alternative," *Synthese* (May, 1983), p. 188.
- 119 See "Epistemic Operators."
- 120 To defend closure within a relevant-alternatives framework, one cannot define the set of relevant alternatives simply as the set of alternatives that need not be known to be false. See Cohen, "How to be a Fallibilist;" Stine, "Skepticism, Relevant Alternatives, and Deductive Closure."
- 121 For a discussion of these issues and a defense of closure within a relevant alternatives framework, see Cohen, "How to be a Fallibilist;" Stine, "Skepticism, Relevant Alternatives, and Deductive Closure."
- 122 In "Knowledge and Context" and "Knowledge, Context, and Social Standards," I argue that the applicability of this view to nonskeptical issues provides an independent motivation for its application to skeptical issues.
- 123 "Pragmatic Dimension of Knowledge," p. 375.
- 124 *Ibid.*, p. 375.
- 125 *Ibid.*, pp. 375-6.
- 126 *Ibid.*, p. 376.
- 127 Some may think that the issue is not whether the operator knows, but whether he knows that he knows, cf. David Sanford, "Knowledge and Relevant Alternatives: Comments on Dretske," *Philosophical Studies*, 40 (October, 1981), pp. 000-00. Although I do not agree, it would still follow that attributions of knowledge, second-order or otherwise, can be affected by the importance of what is known.
- 128 *Knowledge and the Flow of Information*, p. 133.
- 129 I first heard an example like this from John Pollock.
- 130 "Pragmatic Dimension of Knowledge," p. 377.
- 131 For discussions of this problem, see Sanford, "Knowledge and Relevant Alternatives;" Dretske, "Pragmatic Dimension of Knowledge;" Goldman, "Discrimination and Perceptual Knowledge."
- 132 See Ernest Sosa, "On Knowledge and Context," *Journal of Philosophy* (October, 1986); Yourgrau, "Knowledge and Relevant Alternatives."
- 133 "Pragmatic Dimension of Knowledge," p. 370.
- 134 Some non-skeptics might have the intuition that all alternatives are relevant, but believe that in everyday instances of knowledge, all alternatives can be, in some sense, eliminated. The important distinction concerning knowledge and those that do not. Whether the latter are such that we can eliminate them or are such that knowledge does not require that they be eliminated is not important in this context. Of course, this issue matters very much for the status of epistemic closure principles. For a discussion of this see Cohen, "How to be a Fallibilist."
- 135 See Barry Stroud, *The Significance of Philosophical Skepticism* (Oxford: Oxford University Press 1974); Unger, *Philosophical Relativity*.

