FRED DRETSKE

THE PRAGMATIC DIMENSION OF KNOWLEDGE

(Received 20 February, 1981)

Knowing that something is so, unlike being wealthy or reasonable, is not a matter of degree. Two people can both be wealthy, yet one be wealthier than the other; both be reasonable, yet one be more reasonable than the other. When talking about people, places and topics (*things* rather than facts), it makes sense to say that one person knows something *better than* another. He knows the city better than we do, knows more Russian history than any of his colleagues, but doesn't know his wife as well as do his friends. But *factual* knowledge, the knowledge *that* something is so, does not admit of such comparisons.[1] If we both know that today is Friday, it makes no sense to say that you know this better than I. A rich man can become richer by acquiring more money, and a person's belief (that today is Saturday, for example) can be made more reasonable by the accumulation of additional evidence, but if a person already knows that today is Friday, there is nothing he can acquire that will make him know it better. Additional evidence will not promote him to a loftier form of knowledge –though it may make him *more certain* of something he already knew. You can boil water beyond its boiling point (e.g., at 300° F) but you are not, thereby, boiling it better. You are simply boiling it at a higher temperature.

In this respect factual knowledge is *absolute*. It is like being pregnant: an all or nothing affair. One person cannot be *more* pregnant, or pregnant *better than* someone else. Those who view knowledge as a form of justified (true) belief typically acknowledge this fact by speaking, not simply of justification, but of *full, complete*, or *adequate* justification. Those qualifications on the sort of justification required to know something constitute an admission that knowledge is, whereas justification is not, an absolute idea. For these qualifiers are meant to reflect the fact that there is a certain threshold of justification that must be equalled or exceeded if knowledge is to be obtained, and *equalling or exceeding this threshold* is, of course, an absolute idea. I can have a better justirication than you, but my justification cannot

Philosophical Studies 40 (1981) 363–378. 0031-8116/81/0403-0363$01.60
Copyright © 1981 by D. Reidel Publishing Co., Dordrecht, Holland, and Boston, U.S.A.

SEP-10-2009 22:14
P.001/010
8/25/00 11:52 AM

be more adequate (more sufficient, more full) than yours. If my justification is complete in the intended sense, then your justification cannot be more complete.

Philosophers who view knowledge as some form of justified true belief are generally reluctant to talk about this implied threshold of justification. Just how much evidence or justification, one wants to ask, is *enough* to qualify as an adequate, a full, or a complete justification? If the level or degree of justification is represented by real numbers between 0 and 1 (indicating the conditional probability of that for which one has evidence or justification), any threshold less than 1 seems arbitrary. Why, for example, should a justification of 0.95 be good enough to know something when a justification of 0.94 is not adequate? And if one can know P because one's justification is 0.95 and know Q because one's justification is similarly high, is one excluded from knowing P *and* Q because the justification for their joint occurrence has (in accordance with the multiplicative rule in probability theory) dropped below 0.95?

Aside, though, from its arbitrariness, any threshold of justification less than 1 seems to be *too low*. For examples can easily be given in which such thresholds are exceeded without the justification being *good enough* (by ordinary intuitive standards) for knowledge. For example, if the threshold is set at 0.95, one need only think of a bag with 96 white balls and 4 black balls in it. If someone draws a ball at random from this bag, the justification for believing it to be white exceeds the 0.95 threshold. Yet, it seems clear (to me at least) that such a justification (for believing that a white ball has been drawn) is *not* good enough. Someone who happened to draw a white ball, and believed they drew a white ball on the basis of this justification, would not know that they drew a white ball.

Examples such as this suggest (though they do not, of course, prove) that the absolute, non-comparative, character of knowledge derives from the absoluteness, or conclusiveness, of the justification required to know. If I know that the Russians invaded Afghanistan, you can't know this better than I know it because in order to know it I must already have an optimal, or conclusive justification (a justification at the level of 1), and you can't do better than that. I have explored this possibility in other papers, and I do not intend to pursue it here.[2] What I want to develop in this paper is a different theme, one that (I hope) helps to illuminate our concept of knowledge by showing how this absolute idea can, despite its absoluteness, remain sensitive to the

shifting interests, concerns and factors influencing its everyday application. In short, I want to explore the way, and the extent to which, this absolute notion exhibits a degree of contextual relativity in its ordinary use.

To do this it will be useful to briefly recapitulate Peter Unger's discussion of absolute concepts.[3] Although he misinterprets its significance, Unger does, I think, locate the important characteristic of this class of concepts. He illustrates the point with the term *flat*. This, he argues, is an absolute term in the sense that a surface is flat only if it is *not at all bumpy or irregular*. Any bumps or irregularities, however small, and insignificant they may be (from a practical point of view), mean that the surface on which they occur is not really flat. It may be *almost* flat, or *very nearly* flat, but (as both these expressions imply) it is not really flat. We do, it seems, compare surfaces with respect to their degree of flatness (e.g., West Texas is flatter than Wisconsin), but Unger argues that this must be understood as a comparison of the degree to which these surfaces approximate flatness. They cannot both be flat and, yet, one be flatter than the other. Hence, if A is flatter than B, then B (perhaps also A) is not really flat. Flatness does not admit of degrees although a surface's nearness to being flat does, and it is this latter magnitude that we are comparing when we speak of one surface being flatter than another.

Unger concludes from this analysis that not many things are really flat. For under powerful enough magnification almost any surface will exhibit *some* irregularities. Hence, contrary to what we commonly say (and, presumably, believe), these surfaces are not really flat. When we describe them as being flat, what we say is literally false. Probably *nothing* is really flat. So be it. This, according to Unger, is the price we pay for having absolute concepts.

If knowledge is absolute in this way, then there should be similar objections to its widespread application to everyday situations. Powerful magnification (i.e., critical inquiry) *should*, and with the help of the skeptic *has*, revealed 'bumps' and 'irregularities' in our evidential posture with respect to most of the things we say we know. There are always, it seems, possibilities that our evidence is powerless to eliminate, possibilities which, until eliminated, block the road to knowledge. For if knowledge, being an absolute concept, requires the elimination of *all* competing possibilities (possibilities that contrast with what is known), then, clearly we seldom, if ever, satisfy the conditions for applying the concept.

This skeptical conclusion is unpalatable to most philosophers. Unger

SEP-10-2009 22:14 P.002/010

endorses it. Knowledge, according to him, is an absolute concept that, like flatness, has very little application to our bumpy, irregular world.

I have in one respect already indicated my agreement with Unger. Knowledge *is* an absolute concept (I disagree with him, however, about the source of this absoluteness; Unger finds it in the *certainty* required for knowledge; I find it in the *justification* required for knowledge). Unlike Unger, though, I do not derive skeptical conclusions from this fact. I will happily admit that *flat* is an absolute concept, and absolute in roughly the way Unger says it is, but I do not think this shows that nothing is really 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. 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.

Similarly, a road can be perfectly flat even though one can *feel* and *see* irregularities in its surface, irregularities which, were they to be found on the surface of, say, a mirror would mean that the mirror's surface was not really flat. Large mice are not large animals and flat roads are not necessarily flat surfaces. The Flat Earth society is certainly an anachronism, but they are not denying the existence of ant hills and gopher holes.

Absolute concepts depict a situation as being completely devoid of a certain sort of thing: *bumps* in the case of flatness and *objects* in the case of emptiness. The fact that there can be *nothing* of this sort present for the concept to be satisfied is what makes it an absolute concept. It is why if X is empty, Y cannot be emptier. Nonetheless, when it comes to determining what *counts* as a thing of this sort (a bump or an object), and hence what counts against a correct application of the concept, we find the criteria or standards peculiarly spongy and relative. What counts as a thing for assessing the emptiness of my pocket may not count as a thing for assessing the emptiness of a park, a warehouse, or a football stadium. Such concepts, we might say, are *relationally absolute*; absolute, yes, but only relative to a certain standard. We might put the point this way: to be empty is to be *devoid of all relevant things*, thereby exhibiting, simultaneously, the absolute (in the world 'all') and relative (in the word 'relevant') character of this concept.

If, as I have suggested, knowledge is an absolute concept, we should expect it to exhibit this kind of *relationally* absolute character. This, indeed, is the possibility I mean to explore in this paper. What I propose to do is to use what I have called relationally absolute concepts as a model for understanding knowledge. In accordance with this approach (and in harmony with an earlier suggestion) I propose to think of knowledge as an evidential state in which *all relevant alternatives* (to what is known) *are eliminated*. This makes knowledge an absolute concept but the restriction to *relevant* alternatives makes it, like *empty* and *flat*, applicable to this epistemically bumpy world we live in.

Why do this? What are the advantages? A partial catalog of benefits follows:

(1) A growing number of philosophers are able to find, or so they claim, a pragmatic, social, or communal dimension to knowledge.[4] A variety of examples indicate, or seem to these philosophers to indicate, that knowledge depends, not *just* on the evidential status of the knower vis-à-vis what is known, but on such factors as the general availability, and proximity, of (misleading) counter-evidence, on the sorts of things that are commonly taken for granted by others in the relevant community, on the interests and purposes of speaker (in claiming to know) and listeners (in being told that someone knows), and the importance or significance of *what* is known or someone's knowing it. I, personally, happen to think that most of these examples show nothing of the kind. These factors affect, not *whether* something is known, but whether it is reasonable to *say* you know or to *think* you know. But, for the moment, I do not want to argue the point. I merely wish to point out that in so far as knowledge *is* a function of such pragmatic, social or communal factors, the present approach to its analysis can absorb this relativity without compromising the absoluteness of knowledge itself. The social or pragmatic dimension to knowledge, if it exists at all, has to do with what *counts* as a relevant alternative, a possibility that must be evidentially excluded, in order to have knowledge. It does not change the fact that to

SEP-10-2009 22:14 P.003/010

know one must be in a position to exclude *all* such possibilities. It does not alter the fact that one must have, in this sense, an optimal justification – one that eliminates every (relevant) possibility of being mistaken.

(2) Secondly, this approach to the analysis of knowledge helps to avoid the proliferation of *senses* that sometimes threatens to engulf epistemological discussions. We don't have different senses of the verb 'to know' – a strong sense here, a weak sense there – but *one* sense with different applications. We don't have two senses of the word 'empty' – one for pockets and one for warehouses. We have one sense (or meaning) with a difference in what counts as a thing.

(3) Thirdly, 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 wilder 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.

(4) Finally, this approach to the analysis of knowledge gives us the kind of machinery we need to handle the otherwise puzzling examples that are becoming more frequent in the epistemological literature. Consider yet one more example (one *more* because this one, I think, combines elements of several of the more familiar examples). An amateur bird watcher spots a duck

on his favorite Wisconsin pond. He quickly notes its familiar silhouette and markings and makes a mental note to tell his friends that he saw a Gadwall, a rather unusual bird in that part of the midwest. Since the Gadwall has a distinctive set of markings (black rump, white patch on the hind edge of the wing, etc.), markings that no other North American duck exhibits, and these markings were all perfectly visible, it seems reasonable enough to say that the bird-watcher *knows* that yonder bird is a Gadwall. He can see that it is.

Nevertheless, a concerned ornithologist is poking around in the vicinity, not far from where our bird-watcher spotted his Gadwall, looking for some trace of Siberian Grebes. Grebes are duck-like water birds, and the Siberian version of this creature is, when it is in the water, very hard to distinguish from a Gadwall duck. Accurate identification requires seeing the birds in flight since the Gadwall has a white belly and the Grebe a red belly-features that are not visible when the birds are in the water. The ornithologist has a hypothesis that some Siberian Grebes have been migrating to the midwest from their home in Siberia, and he and his research assistants are combing the midwest in search of confirmation.

Once we embellish our simple story in this way, intuitions start to diverge on whether our amateur bird-watcher does indeed know that yonder bird is a Gadwall duck (we are assuming, of course, that it *is* a Gadwall). Most people (I assume) would say that he did *not* know the bird to be a Gadwall if there actually were Siberian Grebes in the vicinity. It certainly sounds strange to suppose that he could give assurances to the ornithologist that the bird he saw was *not* a Siberian Grebe (since he knew it to be a Gadwall duck). But what if the ornithologist's suspicions are unfounded. None of the Grebes have migrated. Does the bird-watcher still not know what he takes himself to know. Is, then, the simple presence of an ornithologist, with his false hypothesis, enough to rob the bird-watcher of his knowledge that the bird on the pond is a Gadwall duck? What if we suppose that the Siberian Grebes, because of certain geographical barriers, *cannot* migrate. Or suppose that there really are no Siberian Grebes – the existence of such a bird being a delusion of a crackpot ornithologist. We may even suppose that, in addition to there being no grebes, there is no ornithologist of the sort I described, but that people in the area believe that there is. Or *some* people believe that there is. Or the bird-watcher's *wife* believes that there is and, as a result, expresses skepticism about his claim to know that what he saw was a Gadwall duck. Or, finally, though no one believes any of this, some of the locals are interested

SEP-10-2009 22:14 P.004/010

in whether or not our birdwatcher *knows* that there are no look-alike migrant grebes in the area.

Somewhere in this progression philosophers, most of them anyway, will dig in their heels and say that the bird-watcher really *does* know that the bird he sees is a Gadwall, and that he knows this despite his inability to justifiably rule out certain alternative possibilities. For example, if there are no look-alike grebes and no ornithologist of the sort I described, but the bird-watcher's wife believes that there are (a rumour she heard from her hairdresser), this does not rob him of his knowledge that the bird he saw as a Gadwall. He needn't be able to rule out the possibility that there are, somewhere in the world, look-alike grebes that have migrated to the midwest in order to know that the bird he saw was a Gadwall duck. These other possibilities are (whether the bird-watcher realizes it or not) simply too remote.

Most philosophers will dig in their heels here because they realize that if they don't, they are on the slippery slope to skepticism with nothing left to hang onto. If false rumours about look-alike grebes and ornithologists can rob an expert bird-watcher of his knowledge that a bird seen in good light, and under ideal conditions, is a Gadwall duck, then similarly false rumours, suspicions or even conjectures about deceptive demons or possible tricks will rob everyone of almost everything they know. 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 bird-watcher 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. 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. In the first case 'empty' implies having no food in it; in the second it implies having no shelves, brackets and hardware in it.

These, then, are some of the advantages to be derived from this approach to the analysis of knowledge. They are, however, advantages that can only be harvested if certain questions can be given reasonable answers: in particular (a) what makes a possibility relevant? (b) If, in order to know, one must rule out all relevant alternatives, how is this 'elimination' to be understood? What does it take, evidentially, to 'rule out' an alternative? (c) Is it possible, as this type of analysis suggests, for one to know something at one time and, later, not know it (due to the introduction of another relevant alternative) without forgetting it? (c) Can one make it easier to know things by remaining ignorant of what are, for others, relevant possibilities?

These, and many more questions, need answers if this framework for the analysis of knowledge is to be anything more than suggestive. Since I cannot here (or anywhere else, for that matter) provide answers to all these questions, I will try, in the time remaining, to fill in some of the large gaps.

Call the *Contrasting Set* (*CS*) the class of situations that are necessarily eliminated by what is known to be the case. That is, if *S* knows that *P*, then *Q* is in the *CS* (of *P*) if and only if, given *P*, necessarily not-*Q*. In our bird-watcher's example, the bird's being a Siberian Grebe (or any kind of grebe at all) is in the *CS* of our bird-watcher's knowledge, or putative knowledge, that it is a Gadwall duck. So is its being an elephant, a hummingbird, a holographic image, or a figment of his imagination. Furthermore, let us call the set of possible alternatives that a person must be in an evidential position to exclude (when he knows that *P*) the *Relevancy Set* (*RS*). In saying that he must be in a position to exclude these possibilities I mean that his evidence or justification for thinking these alternatives are *not* the case must be good enough to say he *knows* they are not the case. Items in the *CS* that are not in the *RS* I shall call irrelevant alternatives. These are items which, though their existence is incompatible with what is known to be the case, the knower *need not* (though he may) have a justification for thinking do not exist. Under normal conditions (the kind of conditions that I assume prevail in the world today) the possibility of something's being a look-alike grebe, though it is a member of the contrasting set, is not a member of the relevancy set of a bird-watcher's knowledge that what he sees is a Gadwall duck (in the kind of circumstances I described).[5] On the other hand, its being an eagle, a Mallard, or a Loon *are* members of the relevancy set since if the bird watcher could not eliminate these possibilities (sufficient unto knowing that it was not an eagle, a Mallard or a loon) on the basis of the bird's appearance and behavior, then he would not know that it was a Gadwall.

What we are suggesting here is that the *RS* is always a proper subset of the *CS* and, moreover, may not be the same *RS* from situation to situation even though what is known remains the same. The situation can be diagrammed as follows:

SEP-10-2009 22:14 P.005/010

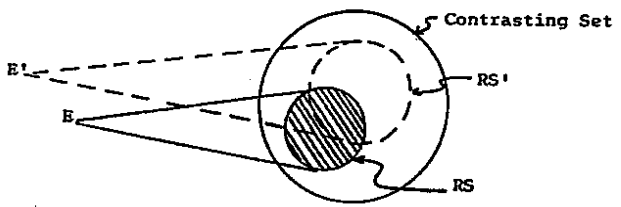

Fig. 1

The solid lines indicate a *RS* and the corresponding piece of evidence that would be required to know with this *RS*. With a different *RS* (*RS'*), indicated by dotted lines, different evidence would be required. If Siberian Grebes are in the Relevancy Set, then additional, more elaborate, evidence is required to know that yonder bird is a Gadwall than in the normal situation. Since the bellies are of different color, one might, for example, be able to tell that it was a Gadwall by watching it in flight. The point, however, is that something more would be needed than was available in the original, normal situation.

In terms of this kind of diagram, a skeptic could be represented as one who took *RS* = *CS* in all cases. One's evidence must be comprehensive enough to eliminate all contrasting possibilities – there being no irrelevant alternatives.

Once the mistake is made of identifying *RS* with *CS* the pressure (on non-skeptics) for lowering the standards of justification (requisite for knowing) becomes irresistible. For if in order to know that *P* one must be jusstified in rejecting *all* members of the *CS* (not just all members of the *RS*), then one can no longer expect very impressive levels of justification for what people know to be the case. If the evidence our bird watcher has for believing the bird to be a Gadwall duck (wing markings, etc.) is also supposed to justify the proposition that it is *not* a look-alike grebe, then, obviously, the justification is nowhere near conclusive. What some philosophers seem inclined to conclude from this is that knowledge does not require conclusive evidence. The reasoning is simple: the bird-watcher knows it is a Gadwall; he doesn't have conclusive reasons (he can't exclude the possibility that it is a look-alike grebe); therefore knowledge does not require conclusive reasons. But

this, I submit, is a fallacy, a misunderstanding of what needs to be conclusively excluded in order to know. Such reasoning is analogous to arguing that to be empty an object can have a few hundred things in it, and to conclude this on the basis of the undeniable fact that empty classrooms, warehouses, and buildings generally have at least a hundred things in them.

But what determines the membership of a relevancy set? A relevancy set, you will recall, is a set of situations each member of which contrasts with what is known to be the case, and must be evidentially excluded if one is to know. Are there criteria for membership in this set? I'm now going to stick my neck out by saying what some of the considerations are that determine the membership of these sets. I do not expect much agreement.

(1) The first point has to do with the way we use contrastive focusing to indicate the range of relevant alternatives. I have discussed this phenomenon in another place, so let me give just one example to illustrate the sort of thing I have in mind.[6] Someone claiming to know that Clyde *sold* his typewriter to Alex is not (necessarily) claiming the same thing as one who claims to know that Clyde sold his typewriter *to Alex*. The sentence we use to express what they know is the same, of course, but they reflect, and are designed to reflect, different relevancy sets. A person who knows that Clyde *sold* his typewriter to Alex must be able to rule out the possibility that he *gave* it to him, or that he *loaned* it to him, or (perhaps) that he merely *pretended* to sell it to him. But he needs only a nominal justification, if he needs any justification at all, for thinking it was Alex to whom he sold it. He has to be right about its *being* Alex, of course, but he isn't claiming to have any special justification for thinking it was Alex rather than, say, his twin brother Albert. On the other hand, the person who knows that Clyde sold his typewriter *to Alex* is claiming to know that it wasn't Albert and is, therefore, expected to be in possession of evidence bearing on the identity of the recipient. But, in this second case, the knower needs only a nominal justification for the belief that Clyde *sold* him the typewriter rather than, say, loaned it to him. He certainly needn't be able to exclude the possibility that the entire transaction was a sham designed to fool the IRS.

(2) A second point, related to the first, is the way the subject term chosen to give verbal expression to what is known often functions to restrict the range of relevant alternatives.[7] Once again, an example will have to suffice. If I say that I could tell that your sister was amused by my funny story, I do not thereby claim to know that she is really your sister, really

SEP-10-2009 22:14 P.006/010

a human being (rather than a cleverly contrived robot), or really the sort of creature that could experience amusement. These possibilities, though certainly relevant to the truth of what I say in the sense that if they were realized I would not know what I say I know are not possibilities that I need be in an evidential position to exclude to know that your sister was amused by my joke. I was, as it were, *taking it for granted* that she was your sister (hence, a human being, a creature that could experience amusement), and I was claiming to know something about the thing so referred to. On the other hand, if I said that I could tell that the object in the corner (that happened to be your sister) was amused by my funny story, the possibility that it is a robot becomes a relevant alternative, one that I am (by this choice of words) accepting epistemic responsibility for excluding.

(3) Thirdly, in saying that we know we often reveal, either directly or indirectly, *how* we know. I could *see* that the tire was flat, could tell (by the way they *behaved*) that they were in love, *heard* them making plans to go, learned (from the *newspapers*) that the auto workers went out on strike, and used *my pocket calculator* to get the answer. The *way* we come to know, the channel (so to speak) over which we obtain our information, is, I submit, always the locus of irrelevant alternatives. Others can challenge the reliability of this channel (our visual system, our auditory system, the newspapers, the pocket calculator), and if it turns out unreliable in some way they will thereby have discredited our claim to knowledge. But others cannot discredit our claim to knowledge merely by pointing out that the channel over which we received our information *could be* unreliable or that we do not *know* it to be reliable. Possible malfunctions in the channel over which we receive information (combined with a resulting false message) are members of the contrasting set but they are not members of the relevancy set. To say that one can see, by the newspapers, that the auto workers are going on strike is to advance a knowledge claim (that the auto workers are going on strike) on the *assumption* of the newspapers reliability. *If* the newspapers *are* a reliable source of such information, then the claimant does know what he says he knows, and he knows it in the way he says he knows it. One cannot undermine this claim by raising possibilities about deceptive newspaper stories or arguing that the claimant does not know that the newspapers, or this newspaper, is reliable. He never said he did know this. What he did say is that he knew the auto workers were going out on strike while simultaneously

disclosing what he was taking for granted which, *if true*, allowed him to know this.

I take the same to be true about our sensory systems when we come to know something by seeing, hearing, tasting and touching. This is the function of our frequent indications (when advancing a knowledge claim) of the manner in which we came to know. We are, by this device, tipping off our listeners, helping them to identify which possibilities are irrelevant to what we are claiming to know.

(4) Fourthly, some people, I am sure, would insist that a pertinent factor influencing the size and membership of the relevancy set is the importance (for speaker and listeners) of what is known or of someone's knowing it. There is a difference between driving by a string of gasoline stations and driving in the middle of the desert. Running out of gas in the first case may be merely an inconvenience; in the latter case it may be a matter of life and death. This makes a difference between knowing (by glancing at your fuel gauge) that you still have some gas in your tank. The implications of being wrong in these two cases are much different – *so* different (some would claim) that additional precautions must be taken (to rule out certain possibilities) in the second case if one is to *know* that one still has some gasoline. And there is even a bigger difference between these cases and knowing that the coolant liquid surrounding the reactor on Three Mile Island is at a safe level by glancing at a similar kind of gauge. 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. 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.

There is, I admit, some appeal to this point, but I think it 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

SEP-10-2009 22:14 P.007/010

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. I admit that the operators *should not* rely on a single gauge, and certainly not one manufactured under such casual quality control, but if they *do* rely on it, I don't see any basis for denying that they know. They should be nervous, but this nervousness is not to be explained by their failure to know what the coolant level is, but by their uncertainty as to when (due to gauge malfunction) they *stop* knowing it.

(5) Finally, we come to the difficult question, the question of when an alternative (not otherwise excluded as irrelevant by one of the considerations already discussed) is just *too remote* to qualify as relevant. In the case of our bird watcher, some philosophers, thinking to turn the tables on the skeptic (by drastically diminishing the relevance set), have suggested that an alternative only becomes relevant when there are positive reasons for thinking it is, or may be, realized. Doubt can also be irrational, and if there are no reasons to doubt, mere possibilities are irrelevant to whether what is believed is known.

This, obviously, is an over-reaction. The Wisconsin lakes could be loaded with migrant Siberian Grebes without the bird watcher having any reason to think that such look-alike birds actually existed. His lack of any reason to doubt, his ignorance of the possibility that what he sees is a grebe and not a Gadwall, is irrelevant. The mere possibility is in this case enough to show he doesn't know.

This shows that having a reason (evidence) to think X is a genuine possibility is not a necessary condition for X's being a relevant alternative. Perhaps, though, it is sufficient. Perhaps, that is, a reasonable (justified) belief that yonder bird *might* be a look-alike grebe (whether or not this belief is true) is enough to make its being a look-alike grebe a relevant possibility.

But if a person really does believe that the bird could be a grebe, aside from the question of whether or not this belief is reasonable, he surely fails to have the kind of belief requisite to knowing it is a Gadwell. He certainly doesn't think he knows it is a Gadwall. I do not know exactly how to express the belief condition on knowledge, but it seems to me that anyone who believes (reasonably or not) that he *might* be wrong fails to meet it.[8] And so the present suggestion is irrelevant to our problem. It describes conditions in which the subject fails to know but only by robbing him of the belief requisite to knowledge.

It may be thought that the mere presence of evidence that one might be wrong, assuming this evidence does not affect one's willingness to believe, is enough to make the respect in which one (according to this evidence) might be wrong a relevant alternative. This has the unfortunate consequence that one can rob a person, indeed a whole community, of its knowledge by spreading a false rumour. I can, for example, tell the bird-watcher that I just met an ornithologist looking for migrant grebes. Once this message is relayed to the bird watcher, even if he rejects it as a silly fabrication, he no longer knows that the bird he saw was a Gadwall duck. And, as a consequence, the whole community is robbed of its knowledge that their local pond was visited by a rather rare bird (a Gadwall duck). The mere fact that I have given them a reason to think that the bird could be a look-alike grebe,[9] whether or not they accept this as a reason, implies that, lacking evidence that it was not a grebe, they do not know it was a Gadwall.

Without dragging the dialectic out any longer, let me simply say what such considerations suggest to me. They suggest that the difference between a relevant and an irrelevant alternative resides, not in what we happen to *regard* as a real possibility (whether reasonably or not), but in the kind of possibilities that actually exist in the objective situation. Whether or not our bird watcher knows that the bird he sees is a Gadwall depends on whether or not, in some objective sense, it could be a look-alike grebe (or any other similar looking creature). If, as a matter of fact, there are no look-alike grebes, that settles the matter. He knows it is a Gadwall. If there are grebes, but due to certain geographical barriers, they are confined to their Siberian habitat, then, once again, the possibility of the bird's being a grebe, though remaining a logical possibility, is not a relevant possibility. They, the grebes, cannot migrate to the midwest.

If, however, there are grebes, and they can migrate, but just have not done so, the case becomes more ambiguous. I think, however, that we now have a genuine possibility, a relevant alternative. By hypothesis the bird-watcher does not know it is not a migrant grebe, and however improbable this possibility may be, there is nothing the bird watcher has (either in the look of the bird or in general background information) that excludes the possibility that what he is looking at is a migrant grebe. He does not, therefore, know it to be

SEP-10-2009 22:14 P.008/010

a Gadwall. He will, no doubt, say he knows. And everyone else may agree and, as a result, think *they* know (having been told by someone who knows). But the truth lies elsewhere. It is, I suggest, tantamount to saying that the bottle is empty when there is a drop left. No one is going to quarrel with this description since all the relevant implications (e.g., we can't make another martini) are true. But the claim itself is false.

University of Wisconsin-Madison

NOTES

[1] I know we sometimes say things that suggest a comparison of this sort (e.g., No one knows better than I that there are a lot of mosquitos in the Northwest Territories), but I take such constructions to be describing, not better knowledge, but more direct, more compelling, kinds of evidence.

[2] 'Conclusive reasons', Australasian Journal of Philosophy (May 1971) and: Seeing and Knowing (University of Chicago Press, 1969).

[3] Peter Unger, 'A defense of skepticism', Philosophical Review 80 (1971).

[4] I have in mind Harman's discussion in: Thought (Princeton, 1973) of evidence one does not possess, Goldman's barn example in 'Discrimination and perceptual knowledge', The Journal of Philosophy 73.20 (1976), the sorts of examples appearing in various Defeasibility analyses of knowledge (see Keith Lehrer and Thomas Paxson, Jr., 'Knowledge: Undefeated justified true belief', Journal of Philosophy 66.8 (1969) and Peter Klein 'A proposed definition of propositional knowledge', Journal of Philosophy, 68.16 (1971)), Ernest Sosa's recommendation (in 'How do you know?', American Philosophical Quarterly (1974), Volume 11, Number 2) that we must depart from the traditional conception of knowledge by putting in relief the relativity of knowledge to an epistemic community (p. 117), and David Annis' 'A contextualist theory of epistemic justification', American Philosophical Quarterly, 15.3 (1978) in which the basic model of justification (and, presumably of knowledge) revolves around a person's being able to meet certain objections. The trend here, if this is a trend, seems to be toward the kind of relativity espoused by Thomas Kuhn in his: The Structure of Scientific Revolutions (Chicago, 1962).

[5] Though there are grebes, and some of them look like ducks, there are (to the best of my knowledge) no Siberian Grebes that look like Gadwall ducks. This part of my story was pure invention.

[6] In 'Contrastive statements', The Philosophical Review (October 1972).

[7] I tried to describe the way this works with perceptual claims in: Seeing and Knowing (Chicago, 1969), pp. 93–112.

[8] We needn't suppose that for S to know that P, S must believe that he can't be wrong. But it does seem reasonable to insist that if S knows that P, he does not believe that he might be wrong. In other words, if the bird-watcher really believes that the bird he sees might be a grebe, then he does not know it is a Gadwall.

[9] I assume here that my saying, "There is an ornithologist in the area looking for migrant grebes, a species that looks very much like a Gadwall duck" is *prima facie* evidence that there *is* an ornothologist in the area looking for migrant grebes. If the bird-watcher ignores me (as we are assuming he does), he nonetheless has been given evidence that the bird he saw might have been a grebe.

8/25/00 11:52 AM
SEP-10-2009 22:14
P.009/010