
Discrimination and Perceptual Knowledge

Author(s): Alvin I. Goldman

Source: *The Journal of Philosophy*, Nov. 18, 1976, Vol. 73, No. 20 (Nov. 18, 1976), pp. 771-791

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.com/stable/2025679>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*

JSTOR

THE JOURNAL OF PHILOSOPHY

VOLUME LXXIII, NO. 20, NOVEMBER 18, 1976

DISCRIMINATION AND PERCEPTUAL KNOWLEDGE *

THIS paper presents a partial analysis of perceptual knowledge, an analysis that will, I hope, lay a foundation for a general theory of knowing. Like an earlier theory I proposed,¹ the envisaged theory would seek to explicate the concept of knowledge by reference to the causal processes that produce (or sustain) belief. Unlike the earlier theory, however, it would abandon the requirement that a knower's belief that *p* be causally connected with the fact, or state of affairs, that *p*.

What kinds of causal processes or mechanisms must be responsible for a belief if that belief is to count as knowledge? They must be mechanisms that are, in an appropriate sense, "reliable." Roughly, a cognitive mechanism or process is reliable if it not only produces true beliefs in actual situations, but would produce true beliefs, or at least inhibit false beliefs, in relevant counterfactual situations. The theory of knowledge I envisage, then, would contain an important counterfactual component.

To be reliable, a cognitive mechanism must enable a person to *discriminate* or *differentiate* between incompatible states of affairs. It must operate in such a way that incompatible states of the world would generate different cognitive responses. Perceptual mechanisms illustrate this clearly. A perceptual mechanism is reliable to the extent that contrary features of the environment (e.g., an object's being red, versus its being yellow) would produce contrary perceptual states of the organism, which would, in turn, produce suitably

* An early version of this paper was read at the 1972 Chapel Hill Colloquium. Later versions were read at the 1973 University of Cincinnati Colloquium, and at a number of other philosophy departments. For comments and criticism, I am especially indebted to Holly Goldman, Bruce Aune, Jaegwon Kim, Louis Loeb, and Kendall Walton.

¹ "A Causal Theory of Knowing," this JOURNAL LXIV, 12 (June 22, 1967): 357-372; reprinted in M. Roth and L. Galis, eds., *Knowing* (New York: Random House, 1970).

different beliefs about the environment. Another belief-governing mechanism is a reasoning mechanism, which, given a set of antecedent beliefs, generates or inhibits various new beliefs. A reasoning mechanism is reliable to the extent that its functional procedures would generate new true beliefs from antecedent true beliefs.

My emphasis on discrimination accords with a sense of the verb 'know' that has been neglected by philosophers. The O.E.D. lists one (early) sense of 'know' as "to *distinguish* (one thing) *from* (another)," as in "I know a hawk from a handsaw" (*Hamlet*) and "We'll teach him to know Turtles from Jayes" (*Merry Wives of Windsor*). Although it no longer has great currency, this sense still survives in such expressions as "I don't know him from Adam," "He doesn't know right from left," and other phrases that readily come to mind. I suspect that this construction is historically important and can be used to shed light on constructions in which 'know' takes propositional objects. I suggest that a person is said to know that p just in case he *distinguishes* or *discriminates* the truth of p from relevant alternatives.

A knowledge attribution imputes to someone the discrimination of a given state of affairs from possible alternatives, but not necessarily all logically possible alternatives. In forming beliefs about the world, we do not normally consider all logical possibilities. And in deciding whether someone knows that p (its truth being assumed), we do not ordinarily require him to discriminate p from all logically possible alternatives. Which alternatives are, or ought to be considered, is a question I shall not fully resolve in this paper, but some new perspectives will be examined. I take up this topic in section I.

I

Consider the following example. Henry is driving in the countryside with his son. For the boy's edification Henry identifies various objects on the landscape as they come into view. "That's a cow," says Henry, "That's a tractor," "That's a silo," "That's a barn," etc. Henry has no doubt about the identity of these objects; in particular, he has no doubt that the last-mentioned object is a barn, which indeed it is. Each of the identified objects has features characteristic of its type. Moreover, each object is fully in view, Henry has excellent eyesight, and he has enough time to look at them reasonably carefully, since there is little traffic to distract him.

Given this information, would we say that Henry *knows* that the object is a barn? Most of us would have little hesitation in saying this, so long as we were not in a certain philosophical frame of mind. Contrast our inclination here with the inclination we would have if

we were given some additional information. Suppose we are told that, unknown to Henry, the district he has just entered is full of papier-mâché facsimiles of barns. These facsimiles look from the road exactly like barns, but are really just façades, without back walls or interiors, quite incapable of being used as barns. They are so cleverly constructed that travelers invariably mistake them for barns. Having just entered the district, Henry has not encountered any facsimiles; the object he sees is a genuine barn. But if the object on that site were a facsimile, Henry would mistake it for a barn. Given this new information, we would be strongly inclined to withdraw the claim that Henry *knows* the object is a barn. How is this change in our assessment to be explained?

Note first that the traditional justified-true-belief account of knowledge is of no help in explaining this change. In both cases Henry truly believes (indeed, is certain) that the object is a barn. Moreover, Henry's "justification" or "evidence" for the proposition that the object is a barn is the same in both cases. Thus, Henry should either know in both cases or not know in both cases. The presence of facsimiles in the district should make no difference to whether or not he knows.

My old causal analysis cannot handle the problem either. Henry's belief that the object is a barn is caused by the presence of the barn; indeed, the causal process is a perceptual one. Nonetheless, we are not prepared to say, in the second version, that Henry knows.

One analysis of propositional knowledge that might handle the problem is Peter Unger's non-accidentality analysis.² According to this theory, *S* knows that *p* if and only if it is not at all accidental that *S* is right about its being the case that *p*. In the initial description of the example, this requirement appears to be satisfied; so we say that Henry knows. When informed about the facsimiles, however, we see that it is accidental that Henry is right about its being a barn. So we withdraw our knowledge attribution. The "non-accidentality" analysis is not very satisfying, however, for the notion of "non-accidentality" itself needs explication. Pending explication, it isn't clear whether it correctly handles all cases.

Another approach to knowledge that might handle our problem is the "indefeasibility" approach.³ On this view, *S* knows that *p*

² "An Analysis of Factual Knowledge," this JOURNAL LXV, 6 (Mar. 21, 1968): 157-170; reprinted in Roth and Galis, *op. cit.*

³ See, for example, Keith Lehrer and Thomas Paxson, Jr., "Knowledge: Undefeated Justified True Belief," this JOURNAL, LXVI, 8 (Apr. 24, 1969): 225-237, and Peter D. Klein, "A Proposed Definition of Propositional Knowledge," *ibid.*, LXVIII, 16 (Aug. 19, 1971): 471-482.

only if S 's true belief is justified *and* this justification is not defeated. In an unrestricted form, an indefeasibility theory would say that S 's justification j for believing that p is defeated if and only if there is some true proposition q such that the conjunction of q and j does not justify S in believing that p . In slightly different terms, S 's justification j is defeated just in case p would no longer be evident for S if q were evident for S . This would handle the barn example, presumably, because the true proposition that there are barn facsimiles in the district is such that, if it were evident for Henry, then it would no longer be evident for him that the object he sees is a barn.

The trouble with the indefeasibility approach is that it is too strong, at least in its unrestricted form. On the foregoing account of "defeat," as Gilbert Harman shows,⁴ it will (almost) always be possible to find a true proposition that defeats S 's justification. Hence, S will never (or seldom) know. What is needed is an appropriate restriction on the notion of "defeat," but I am not aware of an appropriate restriction that has been formulated thus far.

The approach to the problem I shall recommend is slightly different. Admittedly, this approach will raise problems analogous to those of the indefeasibility theory, problems which will not be fully resolved here. Nevertheless, I believe this approach is fundamentally on the right track.

What, then, is my proposed treatment of the barn example? A person knows that p , I suggest, only if the actual state of affairs in which p is true is *distinguishable* or *discriminable* by him from a relevant possible state of affairs in which p is false. If there is a relevant possible state of affairs in which p is false and which is indistinguishable by him from the actual state of affairs, then he fails to know that p . In the original description of the barn case there is no hint of any relevant possible state of affairs in which the object in question is not a barn but is indistinguishable (by Henry) from the actual state of affairs. Hence, we are initially inclined to say that Henry knows. The information about the facsimiles, however, introduces such a relevant state of affairs. Given that the district Henry has entered is full of barn facsimiles, there is a relevant alternative hypothesis about the object, viz., that it is a facsimile. Since, by assumption, a state of affairs in which such a hypothesis holds is indistinguishable by Henry from the actual state of affairs (from his vantage point on the road), this hypothesis is not "ruled out" or "precluded" by the factors that prompt Henry's

⁴ *Thought* (Princeton, N.J.: University Press, 1973), p. 152.

belief. So, once apprised of the facsimiles in the district, we are inclined to deny that Henry knows.

Let us be clear about the bearing of the facsimiles on the case. The presence of the facsimiles does not “create” the possibility that the object Henry sees is a facsimile. Even if there were no facsimiles in the district, it would be possible that the object on that site is a facsimile. What the presence of the facsimiles does is make this possibility *relevant*; or it makes us *consider* it relevant.

The qualifier ‘relevant’ plays an important role in my view. If knowledge required the elimination of all logically possible alternatives, there would be no knowledge (at least of contingent truths). If only *relevant* alternatives need to be precluded, however, the scope of knowledge could be substantial. This depends, of course, on which alternatives are relevant.

The issue at hand is directly pertinent to the dispute—at least one dispute—between skeptics and their opponents. In challenging a claim to knowledge (or certainty), a typical move of the skeptic is to adduce an unusual alternative hypothesis that the putative knower is unable to preclude: an alternative compatible with his “data.” In the skeptical stage of his argument, Descartes says that he is unable to preclude the hypothesis that, instead of being seated by the fire, he is asleep in his bed and dreaming, or the hypothesis that an evil and powerful demon is making it appear to him as if he is seated by the fire. Similarly, Bertrand Russell points out that, given any claim about the past, we can adduce the “skeptical hypothesis” that the world sprang into being five minutes ago, exactly as it then was, with a population that “remembered” a wholly unreal past.⁵

One reply open to the skeptic’s opponent is that these skeptical hypotheses are just “idle” hypotheses, and that a person can know a proposition even if there are “idle” alternatives he cannot preclude. The problem, of course, is to specify when an alternative is “idle” and when it is “serious” (“relevant”). Consider Henry once again. Should we say that the possibility of a facsimile before him is a serious or relevant possibility if there are no facsimiles in Henry’s district, but only in Sweden? Or if a single such facsimile once existed in Sweden, but none exist now?

There are two views one might take on this general problem. The first view is that there is a “correct” answer, in any given situation, as to which alternatives are relevant. Given a complete specification of Henry’s situation, a unique set of relevant alternatives is deter-

⁵ *The Analysis of Mind* (London: Allen & Unwin, 1921), pp. 159–160.

mined: either a set to which the facsimile alternative belongs or one to which it doesn't belong. According to this view, the semantic content of 'know' contains (implicit) rules that map any putative knower's circumstances into a set of relevant alternatives. An analysis of 'know' is incomplete unless it specifies these rules. The correct specification will favor either the skeptic or the skeptic's opponent.

The second view denies that a putative knower's circumstances uniquely determine a set of relevant alternatives. At any rate, it denies that the semantic content of 'know' contains rules that map a set of circumstances into a single set of relevant alternatives. According to this second view, the verb 'know' is simply not so semantically determinate.

The second view need not deny that there are *regularities* governing the alternative hypotheses a speaker (i.e., an attributer or denier of knowledge) thinks of, and deems relevant. But these regularities are not part of the semantic content of 'know'. The putative knower's circumstances do not *mandate* a unique selection of alternatives; but psychological regularities govern which set of alternatives are in fact selected. In terms of these regularities (together with the semantic content of 'know'), we can explain the observed use of the term.

It is clear that some of these regularities pertain to the (description of the) putative knower's circumstances. One regularity might be that the more *likely* it is, given the circumstances, that a particular alternative would obtain (rather than the actual state of affairs), the more probable it is that a speaker will regard this alternative as relevant. Or, the more *similar* the situation in which the alternative obtains to the actual situation, the more probable it is that a speaker will regard this alternative as relevant. It is not only the circumstances of the putative knower's situation, however, that influence the choice of alternatives. The speaker's own linguistic and psychological context are also important. If the speaker is in a class where Descartes's evil demon has just been discussed, or Russell's five-minute-old-world hypothesis, he may think of alternatives he would not otherwise think of and will perhaps treat them seriously. This sort of regularity is entirely ignored by the first view.

What I am calling the "second" view might have two variants. The first variant can be imbedded in Robert Stalnaker's framework for pragmatics.⁶ In this framework, a proposition is a function from

⁶"Pragmatics," in Donald Davidson and Harman, eds., *Semantics of Natural Language* (Boston: Reidel, 1972).

possible words into truth values; the determinants of a proposition are a sentence and a (linguistic) context. An important contextual element is what the utterer of a sentence presupposes, or takes for granted. According to the first variant of the second view, a sentence of the form '*S* knows that *p*' does not determine a unique proposition. Rather, a proposition is determined by such a sentence together with the speaker's presuppositions concerning the relevant alternatives.⁷ Sceptics and nonsceptics might make different presuppositions (both presuppositions being "legitimate"), and, if so, they are simply asserting or denying different propositions.

One trouble with this variant is its apparent implication that, if a speaker utters a knowledge sentence without presupposing a fully determinate set of alternatives, he does not assert or deny any proposition. That seems too strong. A second variant of the second view, then, is that sentences of the form '*S* knows that *p*' express vague or indeterminate propositions (if they express "propositions" at all), which can, but need not, be made more determinate by full specification of the alternatives. A person who *assents* to a knowledge sentence says that *S* discriminates the truth of *p* from relevant alternatives; but he may not have a distinct set of alternatives in mind. (Similarly, according to Paul Ziff, a person who says something is "good" says that it answers to *certain* interests;⁸ but he may not have a distinct set of interests in mind.) Someone who *denies* a knowledge sentence more commonly has one or more alternatives in mind as relevant, because his denial may stem from a particular alternative *S* cannot rule out. But even the denier of a knowledge sentence need not have a full set of relevant alternatives in mind.

I am attracted by the second view under discussion, especially its second variant. In the remainder of the paper, however, I shall be officially neutral. In other words, I shall not try to settle the question of whether the semantic content of 'know' contains rules that map the putative knower's situation into a unique set of relevant alternatives. I leave open the question of whether there is a "correct" set of relevant alternatives, and if so, what it is. To this extent, I also leave open the question of whether sceptics or their opponents are "right." In defending my analysis of 'perceptually knows', however, I shall have to discuss particular examples. In treating these examples I shall assume some (psycho-

⁷ Something like this is suggested by Fred Dretske, in "Epistemic Operators," this JOURNAL, LXVII, 24 (Dec. 24, 1970): 1007-1023, p. 1022.

⁸ That 'good' means *answers to certain interests* is claimed by Ziff in *Semantic Analysis* (Ithaca, N.Y.: Cornell, 1960), ch. vi.

logical) regularities concerning the selection of alternatives. Among these regularities is the fact that speakers do not *ordinarily* think of "radical" alternatives, but are caused to think of such alternatives, and take them seriously, if the putative knower's circumstances call attention to them. Since I assume that radical or unusual alternatives are not *ordinarily* entertained or taken seriously, I may appear to side with the opponents of skepticism. My official analysis, however, is neutral on the issue of skepticism.

II

I turn now to the analysis of 'perceptually knows'. Suppose that Sam spots Judy on the street and correctly identifies her as Judy, i.e., believes she is Judy. Suppose further that Judy has an identical twin, Trudy, and the possibility of the person's being Trudy (rather than Judy) is a relevant alternative. Under what circumstances would we say that Sam *knows* it is Judy?

If Sam regularly identifies Judy as Judy and Trudy as Trudy, he apparently has some (visual) way of discriminating between them (though he may not know how he does it, i.e., what cues he uses). If he does have a way of discriminating between them, which he uses on the occasion in question, we would say that he *knows* it is Judy. But if Sam frequently mistakes Judy for Trudy, and Trudy for Judy, he presumably does not have a way of discriminating between them. For example, he may not have sufficiently distinct (visual) memory "schemata" of Judy and Trudy. So that, on a particular occasion, sensory stimulation from either Judy *or* Trudy would elicit a Judy-identification from him. If he happens to be right that it is Judy, this is just accidental. He doesn't *know* it is Judy.

The crucial question in assessing a knowledge attribution, then, appears to be the truth value of a counterfactual (or set of counterfactuals). Where Sam correctly identifies Judy as Judy, the crucial counterfactual is: "If the person before Sam were Trudy (rather than Judy), Sam would believe her to be Judy." If this counterfactual is true, Sam doesn't know it is Judy. If this counterfactual is false (and all other counterfactuals involving relevant alternatives are also false), then Sam may know it is Judy.

This suggests the following analysis of (noninferential) perceptual knowledge.

S (noninferentially) *perceptually knows that p* if and only if

- (1) *S* (noninferentially) perceptually believes that *p*,
- (2) *p* is true, and
- (3) there is no relevant contrary *q* of *p* such that, if *q* were true (rather than *p*), then *S* would (still) believe that *p*.

Restricting attention to relevant possibilities, these conditions assert in effect that the only situation in which *S* would believe that *p* is a situation in which *p* is true. In other words, *S*'s believing that *p* is sufficient for the truth of *p*. This is essentially the analysis of noninferential knowledge proposed by D. M. Armstrong in *A Materialist Theory of the Mind* (though without any restriction to "relevant" alternatives), and refined and expanded in *Belief, Truth, and Knowledge*.⁹

This analysis is too restrictive. Suppose Oscar is standing in an open field containing Dack the dachshund. Oscar sees Dack and (noninferentially) forms a belief in (P):

(P) The object over there is a dog.

Now suppose that (Q):

(Q) The object over there is a wolf.

is a relevant alternative to (P) (because wolves are frequenters of this field). Further suppose that Oscar has a tendency to mistake wolves for dogs (he confuses them with malamutes, or German shepherds). Then if the object Oscar saw were Wiley the wolf, rather than Dack the dachshund, Oscar would (still) believe (P). This means that Oscar fails to satisfy the proposed analysis with respect to (P), since (3) is violated. But surely it is wrong to deny—for the indicated reasons—that Oscar *knows* (P) to be true. The mere fact that he would erroneously take a wolf to be a dog hardly shows that he doesn't know a *dachshund* to be a dog! Similarly, if someone looks at a huge redwood and correctly believes it to be a tree, he is not disqualified from knowing it to be a tree merely because there is a very small plant he would wrongly believe to be a tree, i.e., a bonsai tree.

The moral can be formulated as follows. If Oscar believes that a dog is present because of a certain way he is "appeared to," then this true belief fails to be knowledge if there is an alternative situation in which a non-dog produces the same belief by means of the same, or a very similar, appearance. But the wolf situation is not such an alternative: although it would produce in him the same belief, it would not be by means of the same (or a similar) appearance. An alternative that disqualifies a true perceptual belief from being perceptual knowledge must be a "perceptual equivalent" of the actual state of affairs.¹⁰ A *perceptual equivalent* of an actual

⁹ *A Materialist Theory of the Mind* (New York: Humanities, 1968), pp. 189 ff., and *Belief, Truth and Knowledge* (New York: Cambridge, 1973), chs. 12 and 13.

¹⁰ My notion of a perceptual equivalent corresponds to Jaakko Hintikka's notion of a "perceptual alternative." See "On the Logic of Perception," in N. S. Care and R. H. Grimm, eds., *Perception and Personal Identity* (Cleveland, Ohio: Case Western Reserve, 1969).

state of affairs is a possible state of affairs that would produce the same, or a sufficiently similar, perceptual experience.

The relation of perceptual equivalence must obviously be relativized to *persons* (or organisms). The presence of Judy and the presence of Trudy might be perceptual equivalents for Sam, but not for the twins' own mother (to whom the twins look quite different). Similarly, perceptual equivalence must be relativized to *times*, since perceptual discriminative capacities can be refined or enhanced with training or experience, and can deteriorate with age or disease.

How shall we specify alternative states of affairs that are candidates for being perceptual equivalents? First, we should specify the *object* involved. (I assume for simplicity that only one object is in question.) As the Judy-Trudy case shows, the object in the alternative state of affairs need not be identical with the actual object. Sometimes, indeed, we may wish to allow non-actual possible objects. Otherwise our framework will be unable in principle to accommodate some of the skeptic's favorite alternatives, e.g., those involving demons. If the reader's ontological sensibility is offended by talk of possible objects, I invite him to replace such talk with any preferred substitute.

Some alternative states of affairs involve the same object but different properties. Where the actual state of affairs involves a certain ball painted blue, an alternative might be chosen involving the same ball painted green. Thus, specification of an alternative requires not only an object, but properties of the object (at the time in question). These should include not only the property in the belief under scrutiny, or one of its contraries, but other properties as well, since the property in the belief (or one of its contraries) might not be sufficiently determinate to indicate what the resultant percept would be like. For full generality, let us choose a *maximal set of* (nonrelational) *properties*. This is a set that would exhaustively characterize an object (at a single time) in some possible world.¹¹

An object plus a maximal set of (nonrelational) properties still does not fully specify a perceptual alternative. Also needed are relations between the object and the perceiver, plus conditions of

¹¹ I have in mind here purely qualitative properties. Properties like *being identical with Judy* would be given by the selected object. If the set of qualitative properties (at a given time) implied which object it was that had these properties, then specification of the object would be redundant, and we could represent states of affairs by ordered pairs of maximal sets of (qualitative) properties and DOE relations. Since this is problematic, however, I include specification of the object as well as the set of (qualitative) properties.

the environment. One relation that can affect the resultant percept is *distance*. Another relational factor is *relative orientation*, both of object vis-à-vis perceiver and perceiver vis-à-vis object. The nature of the percept depends, for example, on which side of the object faces the perceiver, and on how the perceiver's bodily organs are oriented, or situated, vis-à-vis the object. Thirdly, the percept is affected by the current state of the *environment*, e.g., the illumination, the presence or absence of intervening objects, and the direction and velocity of the wind.

To cover all such elements, I introduce the notion of a *distance-orientation-environment* relation, for short, a *DOE relation*. Each such relation is a conjunction of relations or properties concerning distance, orientation, and environmental conditions. One DOE relation is expressed by the predicate 'x is 20 feet from y, the front side of y is facing x, the eyes of x are open and focused in y's direction, no opaque object is interposed between x and y, and y is in moonlight'.

Since the health of sensory organs can affect percepts, it might be argued that this should be included in these relations, thereby opening the condition of these organs to counterfactualization. For simplicity I neglect this complication. This does not mean that I don't regard the condition of sensory organs as open to counterfactualization. I merely omit explicit incorporation of this factor into our exposition.

We can now give more precision to our treatment of perceptual equivalents. Perceptual states of affairs will be specified by ordered triples, each consisting of (1) an object, (2) a maximal set of non-relational properties, and (3) a DOE relation. If *S* perceives object *b* at *t* and if *b* has all the properties in a maximal set *J* and bears DOE relation *R* to *S* at *t*, then the actual state of affairs pertaining to this perceptual episode is represented by the ordered triple $\langle b, J, R \rangle$. An alternative state of affairs is represented by an ordered triple $\langle c, K, R^* \rangle$, which may (but need not) differ from $\langle b, J, R \rangle$ with respect to one or more of its elements.

Under what conditions is an alternative $\langle c, K, R^* \rangle$ a perceptual equivalent of $\langle b, J, R \rangle$ for person *S* at time *t*? I said that a perceptual equivalent is a state of affairs that would produce "the same, or a very similar" perceptual experience. That is not very committal. Must a perceptual equivalent produce exactly the same percept? Given our intended use of perceptual equivalence in the analysis of perceptual knowledge, the answer is clearly No. Suppose that a Trudy-produced percept would be qualitatively distinct from Sam's Judy-produced percept, but similar enough for Sam to mistake

Trudy for Judy. This is sufficient grounds for saying that Sam fails to have knowledge. Qualitative identity of percepts, then, is too strong a requirement for perceptual equivalence.

How should the requirement be weakened? We must not weaken it too much, for the wolf alternative might then be a perceptual equivalent of the dachshund state of affairs. This would have the unwanted consequence that Oscar doesn't know Dack to be a dog.

The solution I propose is this. If the percept produced by the alternative state of affairs would not differ from the actual percept in any respect that is causally relevant to *S*'s belief, this alternative situation is a perceptual equivalent for *S* of the actual situation. Suppose that a Trudy-produced percept would differ from Sam's Judy-produced percept to the extent of having a different eyebrow configuration. (A difference in shape between Judy's and Trudy's eyebrows does not ensure that Sam's percepts would "register" this difference. I assume, however, that the eyebrow difference would be registered in Sam's percepts.) But suppose that Sam's visual "concept" of Judy does not include a feature that reflects this contrast. His Judy-concept includes an "eyebrow feature" only in the sense that the absence of eyebrows would inhibit a Judy-classification. It does not include a more determinate eyebrow feature, though: Sam hasn't learned to associate Judy with distinctively shaped eyebrows. Hence, the distinctive "eyebrow shape" of his actual (Judy-produced) percept is not one of the percept-features that is causally responsible for his believing Judy to be present. Assuming that a Trudy-produced percept would not differ from his actual percept in any *other* causally relevant way, the hypothetical Trudy-situation is a perceptual equivalent of the actual Judy-situation.

Consider now the dachshund-wolf case. The hypothetical percept produced by a wolf would differ from Oscar's actual percept of the dachshund in respects that *are* causally relevant to Oscar's judgment that a dog is present. Let me elaborate. There are various kinds of objects, rather different in shape, size, color, and texture, that would be classified by Oscar as a dog. He has a number of visual "schemata," we might say, each with a distinctive set of features, such that any percept that "matches" or "fits" one of these schemata would elicit a "dog" classification. (I think of a schema not as a "template," but as a set of more-or-less abstract—though iconic—features.¹²) Now, although a dachshund and a wolf would each

¹² For a discussion of iconic schemata, see Michael I. Posner, *Cognition: An Introduction* (Glenview, Ill.: Scott, Foresman, 1973), ch. 3.

produce a dog-belief in Oscar, the percepts produced by these respective stimuli would differ in respects that are causally relevant to Oscar's forming a dog-belief. Since Oscar's dachshund-schema includes such features as having an elongated, sausage-like shape, a smallish size, and droopy ears, these features of the percept are all causally relevant, when a dachshund is present, to Oscar's believing that a dog is present. Since a hypothetical wolf-produced percept would differ in these respects from Oscar's dachshund-produced percept, the hypothetical wolf state of affairs is not a perceptual equivalent of the dachshund state of affairs for Oscar.

The foregoing approach requires us to relativize perceptual equivalence once again, this time to the belief in question, or the property believed to be exemplified. The Trudy-situation is a perceptual equivalent for Sam of the Judy-situation *relative to the property of being* (identical with) *Judy*. The wolf-situation is not a perceptual equivalent for Oscar of the dachshund-situation *relative to the property of being a dog*.

I now propose the following definition of perceptual equivalence:

If object *b* has the maximal set of properties *J* and is in DOE relation *R* to *S* at *t*, if *S* has some percept *P* at *t* that is perceptually caused by *b*'s having *J* and being in *R* to *S* at *t*, and if *P* noninferentially causes *S* to believe (or sustains *S* in believing) of object *b* that it has property *F*, then

$\langle c, K, R^* \rangle$ is a perceptual equivalent of $\langle b, J, R \rangle$ for *S* at *t* relative to property *F* if and only if

- (1) if at *t* object *c* had *K* and were in *R** to *S*, then this would perceptually cause *S* to have some percept *P** at *t*,
- (2) *P** would cause *S* noninferentially to believe (or sustain *S* in believing) of object *c* that it has *F*, and
- (3) *P** would not differ from *P* in any respect that is causally relevant to *S*'s *F*-belief.

Since I shall analyze the *de re*, relational, or transparent sense of 'perceptually knows', I shall want to employ, in my analysis, the *de re* sense of 'believe'. This is why such phrases as 'believe. . . of object *b*' occur in the definition of perceptual equivalence. For present purposes, I take for granted the notion of (perceptual) *de re* belief. I assume, however, that the object of which a person perceptually believes a property to hold is the object he perceives, i.e., the object that "perceptually causes" the percept that elicits the belief. The notion of *perceptual causation* is another notion I take for granted. A person's percept is obviously caused by many objects (or events), not all of which the person is said to perceive. One problem for the theory of perception is to explicate the notion of

perceptual causation, that is, to explain which of the causes of a percept a person is said to perceive. I set this problem aside here.¹³ A third notion I take for granted is the notion of a (noninferential) *perceptual belief*, or perceptual “taking.” Not all beliefs that are noninferentially caused by a percept can be considered perceptual “takings”; “indirectly” caused beliefs would not be so considered. But I make no attempt to delineate the requisite causal relation.

Several other comments on the definition of perceptual equivalence are in order. Notice that the definition is silent on whether J or K contains property F , i.e., whether F is exemplified in either the actual or the alternative states of affairs. The relativization to F (in the definiendum) implies that an F -belief is produced in both situations, not that F is exemplified (in either or both situations). In applying the definition to cases of putative knowledge, we shall focus on cases where F belongs to J (so S 's belief is true in the actual situation) but does not belong to K (so S 's belief is false in the counterfactual situation). But the definition of perceptual equivalence is silent on these matters.

Though the definition does not say so, I assume it is possible for object c to have all properties in K , and possible for c to be in R^* to S while having all properties in K . I do not want condition 1 to be vacuously true, simply by having an impossible antecedent.

It might seem as if the antecedent of (1) should include a further conjunct, expressing the supposition that object b is absent. This might seem necessary to handle cases in which, if c were in R^* to S , but b remained in its actual relation R to S , then b would “block” S 's access to c . (For example, b might be an orange balloon floating over the horizon, and c might be the moon.) This can be handled by the definition as it stands, by construing R^* , where necessary, as including the absence of object b from the perceptual scene. (One cannot *in general* hypothesize that b is absent, for we want to allow object c to be identical with b .)

The definition implies that there is no temporal gap between each object's having its indicated properties and DOE relation and the occurrence of the corresponding percept. This simplification is introduced because no general requirement can be laid down about how long it takes for the stimulus energy to reach the perceiver. The intervals in the actual and alternative states may differ because the stimuli might be at different distances from the perceiver.

III

It is time to turn to the analysis of perceptual knowledge, for which the definition of perceptual equivalence paves the way. I restrict

¹³ I take this problem up in “Perceptual Objects,” forthcoming in *Synthese*.

my attention to perceptual knowledge of the possession, by physical objects, of nonrelational properties. I also restrict the analysis to *noninferential* perceptual knowledge. This frees me from the complex issues introduced by inference, which require separate treatment.

It may be contended that all perceptual judgment is based on inference and, hence, that the proposed restriction reduces the scope of the analysis to nil. Two replies are in order. First, although cognitive psychology establishes that percepts are affected by cognitive factors, such as "expectancies," it is by no means evident that these causal processes should be construed as inferences. Second, even if we were to grant that there is in fact no noninferential perceptual belief, it would still be of epistemological importance to determine whether noninferential perceptual knowledge of the physical world is conceptually possible. This could be explored by considering merely possible cases of noninferential perceptual belief, and seeing whether, under suitable conditions, such belief would count as knowledge.

With these points in mind, we may propose the following (tentative) analysis:

At t S noninferentially perceptually knows of object b that it has property F if and only if

- (1) for some maximal set of nonrelational properties *J* and some DOE relation *R*, object *b* has (all the members of) *J* at *t* and is in *R* to *S* at *t*,
- (2) *F* belongs to *J*,
- (3) (A) *b*'s having *J* and being in *R* to *S* at *t* perceptually causes *S* at *t* to have some percept *P*,¹⁴

¹⁴ Should (3A) be construed as implying that *every* property in *J* is a (perceptual) cause of *P*? No. Many of *b*'s properties are exemplified in its interior or at its backside. These are not causally relevant, at least in visual perception. (3A) must therefore be construed as saying that *P* is (perceptually) caused by *b*'s having (jointly) *all* the members of *J*, and leaving open which, among these members, are individually causally relevant. It follows, however, that (3A) does not require that *b*'s-having-*F*, in particular, is a (perceptual) cause of *P*, and this omission might be regarded as objectionable. "Surely," it will be argued, "*S* perceptually knows *b* to have *F* only if *b*'s-having-*F* (perceptually) causes the percept." The reason I omit this requirement is the following. Suppose *F* is the property of being a dog. Can we say that *b*'s-being-a-dog is a cause of certain light waves' being reflected? This is very dubious. It is the molecular properties of the surface of the animal that are causally responsible for this transmission of light, and hence for the percept.

One might say that, even if the percept needn't be (perceptually) caused by *b*'s-having-*F*, it must at least be caused by microstructural properties of *b* that *ensure* *b*'s-having-*F*. As the dog example again illustrates, however, this is too strong. The surface properties of the dog that reflect the light waves do not *ensure* that the object is a dog, either logically or nomologically. Something could have that surface (on one side) and still have a non-dog interior and backside.

- (B) P noninferentially causes S at t to believe (or sustains S in believing) of object b that it has property F , and
- (C) there is no alternative state of affairs $\langle c, K, R^* \rangle$ such that
- (i) $\langle c, K, R^* \rangle$ is a relevant perceptual equivalent of $\langle b, J, R \rangle$ for S at t relative to property F , and
 - (ii) F does not belong to K .

Conditions 1 and 2 jointly entail the truth condition for knowledge: S knows b to have F (at t) only if b does have F (at t). Condition 3B contains the belief condition for knowledge, restricted, of course, to (noninferential) perceptual belief. The main work of the conditions is done by 3C. It requires that there be no relevant alternative that is (i) a perceptual equivalent to the actual state of affairs relative to property F , and (ii) a state of affairs in which the appropriate object lacks F (and hence S 's F -belief is false).

How does this analysis relate to my theme of a "reliable discriminative mechanism"? A perceptual cognizer may be thought of as a two-part mechanism. The first part constructs percepts (a special class of internal states) from receptor stimulation. The second part operates on percepts to produce beliefs. Now, in order for the conditions of the analysis to be satisfied, each part of the mechanism must be sufficiently discriminating, or "finely tuned." If the first part is not sufficiently discriminating, patterns of receptor stimulation from quite different sources would result in the same (or very similar) percepts, percepts that would generate the same beliefs. If the second part is not sufficiently discriminating, then even if different percepts are constructed by the first part, the same beliefs will be generated by the second part. To be sure, even an indiscriminating bipartite mechanism may produce a belief that, luckily, is true; but there will be other, counterfactual, situations in which such a belief would be false. In this sense, such a mechanism is unreliable. What our analysis says is that S has perceptual knowledge if and only if not only does his perceptual mechanism produce true belief, but there are no relevant counterfactual situations in which the same belief would be produced via an equivalent percept and in which the belief would be false.

Let me now illustrate how the analysis is to be applied to the barn example, where there are facsimiles in Henry's district. Let S =Henry, b =the barn Henry actually sees, and F =the property of

The problem should be solved, I think, by reliance on whether there are relevant perceptual equivalents. If there are no relevant perceptual equivalents in which K excludes being a dog, then the properties of the actual object that are causally responsible for the percept suffice to yield knowledge. We need not require either that the percept be (perceptually) caused by b 's-having- F , nor by any subset of J that "ensures" b 's-having- F .

being a barn. Conditions 1 through 3B are met by letting J take as its value the set of all nonrelational properties actually possessed by the barn at t , R take as its value the actual DOE relation the barn bears to Henry at t , and P take as its value the actual (visual) percept caused by the barn. Condition 3C is violated, however. There *is* a relevant triple that meets subclauses (i) and (ii), i.e., the triple where c =a suitable barn facsimile, K =a suitable set of properties (excluding, of course, the property of being a barn), and R^* =approximately the same DOE relation as the actual one. Thus, Henry does not (noninferentially) perceptually *know* of the barn that it has the property of being a barn.

In the dachshund-wolf case, S =Oscar, b =Dack the dachshund, and F =being a dog. The first several conditions are again met. Is 3C met as well? There is a relevant alternative state of affairs in which Wiley the wolf is believed by Oscar to be a dog, but lacks that property. This state of affairs doesn't violate 3C, however, since it isn't a *perceptual equivalent* of the actual situation relative to being a dog. So this alternative doesn't disqualify Oscar from knowing Dack to be a dog.

Is there another alternative that *is* a perceptual equivalent of the actual situation (relative to being a dog)? We can imagine a DOE relation in which fancy devices between Wiley and Oscar distort the light coming from Wiley and produce in Oscar a Dack-like visual percept. The question here, however, is whether this perceptual equivalent is *relevant*. Relevance is determined not only by the hypothetical object and its properties, but also by the DOE relation. Since the indicated DOE relation is highly unusual, this will count (at least for a nonskeptic) against the alternative's being relevant and against its disqualifying Oscar from knowing.¹⁵

The following "Gettierized" example, suggested by Marshall Swain, might appear to present difficulties. In a dark room there is a

¹⁵ It is the "unusualness" of the DOE relation that inclines us not to count the alternative as relevant; it is not the mere fact that the DOE relation differs from the actual one. In general, our analysis allows knowledge to be defeated or disqualified by alternative situations in which the DOE relation differs from the DOE relation in the actual state of affairs. Our analysis differs in this respect from Fred Dretske's analysis in "Conclusive Reasons," *Australasian Journal of Philosophy*, II, 1 (May 1971): 1-22. Dretske's analysis, which ours resembles on a number of points, considers only those counterfactual situations in which everything that is "logically and causally independent of the state of affairs expressed by P " (7/8) is the same as in the actual situation. (P is the content of S 's belief.) This implies that the actual DOE relation cannot be counterfactualized, but must be held fixed. (It may also imply—depending what P is—that one cannot counterfactualize the perceived object nor the full set of properties J .) This unduly narrows the class of admissible alternatives. Many *relevant* alternatives, that do disqualify knowledge, involve DOE relations that differ from the actual DOE relation.

candle several yards ahead of *S* which *S* sees and believes to be ahead of him. But he sees the candle only indirectly, via a system of mirrors (of which he is unaware) that make it appear as if he were seeing it directly.¹⁶ We would surely deny that *S* knows the candle to be ahead of him. (This case does not really fit our intended analysandum, since the believed property *F* is relational. This detail can be ignored, however.) Why? If we say, with Harman, that all perceptual belief is based on inference, we can maintain that *S* infers that the candle is ahead of him from the premise that he sees whatever he sees *directly*. This premise being false, *S*'s knowing is disqualified on familiar grounds.

My theory suggests another explanation, which makes no unnecessary appeal to inference. We deny that *S* knows, I suggest, because the system of mirrors draws our attention to a perceptual equivalent in which the candle is *not* ahead of *S*, i.e., a state of affairs where the candle is behind *S* but reflected in a system of mirrors so that it appears to be ahead of him. Since the actual state of affairs involves a system of reflecting mirrors, we are impelled to count this alternative as relevant, and hence to deny that *S* knows.

Even in ordinary cases, of course, where *S* sees a candle directly, the possibility of reflecting mirrors constitutes a perceptual equivalent. In the ordinary case, however, we would not count this as relevant; we would not regard it as a "serious" possibility. The Gettierized case impels us to take it seriously because there the actual state of affairs involves a devious system of reflecting mirrors. So we have an explanation of why people are credited with knowing in ordinary perceptual cases but not in the Gettierized case.

The following is a more serious difficulty for our analysis. *S* truly believes something to be a tree, but there is a relevant alternative in which an electrode stimulating *S*'s optic nerve would produce an equivalent percept, which would elicit the same belief. Since this is assumed to be a relevant alternative, it ought to disqualify *S* from knowing. But it doesn't satisfy our definition of a perceptual equivalent, first because the electrode would not be a perceptual cause of the percept (we would not say that *S* *perceives* the electrode), and second because *S* would not believe *of the electrode* (nor *of anything else*) that it is a tree. A similar problem arises where the alternative state of affairs would involve *S*'s having a hallucination.

To deal with these cases, we could revise our analysis of perceptual knowledge as follows. (A similar revision in the definition of percep-

¹⁶ Harman has a similar case, in *Thought*, pp. 22–23. In that case, however, *S* does not see the candle; it is not a cause of his percept. Given our causal requirement for perceptual knowledge, that case is easily handled.

tual equivalence would do the job equally well.) We could reformulate 3C to say that there must neither be a relevant perceptual equivalent of the indicated sort (using our present definition of perceptual equivalence) *nor* a relevant alternative situation in which an equivalent percept occurs and prompts a *de dicto* belief that something has *F*, but where there is nothing that *perceptually* causes this percept and nothing *of which F* is believed to hold. In other words, knowledge can be disqualified by relevant alternative situations where *S* doesn't perceive anything and doesn't have any *de re* (*F*-) belief at all. I am inclined to adopt this solution, but will not actually make this addition to the analysis.

Another difficulty for the analysis is this. Suppose Sam's "schemata" of Judy and Trudy have hitherto been indistinct, so Judy-caused percepts sometimes elicit Judy-beliefs and sometimes Trudy-beliefs, and similarly for Trudy-caused percepts. Today Sam falls down and hits his head. As a consequence a new feature is "added" to his Judy-schema, a mole-associated feature. From now on he will believe someone to be Judy only if he has the sort of percept that would be caused by a Judy-like person with a mole over the left eye. Sam is unaware that this change has taken place and will remain unaware of it, since he isn't conscious of the cues he uses. Until today, neither Judy nor Trudy has had a left-eyebrow mole; but today Judy happens to develop such a mole. Thus, from now on Sam can discriminate Judy from Trudy. Does this mean that he will *know* Judy to be Judy when he correctly identifies her? I am doubtful.

A possible explanation of Sam's not knowing (on future occasions) is that Trudy-with-a-mole is a relevant perceptual equivalent of Judy. This is not Trudy's actual condition, of course, but it might be deemed a relevant possibility. I believe, however, that the mole case calls for a further restriction, one concerning the *genesis* of a person's propensity to form a certain belief as a result of a certain percept. A merely fortuitous or accidental genesis is not enough to support knowledge. I do not know exactly what requirement to impose on the genesis of such a propensity. The mole case intimates that the genesis should involve certain "experience" with objects, but this may be too narrow. I content myself with a very vague addition to our previous conditions, which completes the analysis:

- (4) *S*'s propensity to form an *F*-belief as a result of percept *P* has an appropriate genesis.

Of course this leaves the problem unresolved. But the best I can do here is identify the problem.

IV

A few words are in order about the intended significance of my analysis. One of its purposes is to provide an alternative to the traditional "Cartesian" perspective in epistemology. The Cartesian view combines a theory of knowledge with a theory of justification. Its theory of knowledge asserts that *S* knows that *p* at *t* only if *S* is (fully, adequately, etc.) justified at *t* in believing that *p*. Its theory of justification says that *S* is justified at *t* in believing that *p* only if either (A) *p* is self-warranting for *S* at *t*, or (B) *p* is (strongly, adequately, etc.) supported or confirmed by propositions each of which is self-warranting for *S* at *t*. Now propositions about the state of the external world at *t* are not self-warranting. Hence, if *S* knows any such proposition *p* at *t*, there must be some other propositions which strongly support *p* and which are self-warranting for *S* at *t*. These must be propositions about *S*'s mental state at *t* and perhaps some obvious necessary truths. A major task of Cartesian epistemology is to show that there is some such set of self-warranting propositions, propositions that support external-world propositions with sufficient strength.

It is impossible to canvass all attempts to fulfill this project; but none have succeeded, and I do not think that any will. One can conclude either that we have no knowledge of the external world or that Cartesian requirements are too demanding. I presuppose the latter conclusion in offering my theory of perceptual knowledge. My theory requires no justification for external-world propositions that derives entirely from self-warranting propositions. It requires only, in effect, that beliefs in the external world be suitably caused, where "suitably" comprehends a process or mechanism that not only produces true belief in the actual situation, but would not produce false belief in relevant counterfactual situations. If one wishes, one can so employ the term 'justification' that belief causation of *this* kind counts as justification. In this sense, of course, my theory does require justification. But this is entirely different from the sort of justification demanded by Cartesianism.

My theory protects the possibility of knowledge by making Cartesian-style justification unnecessary. But it leaves a door open to skepticism by its stance on relevant alternatives. This is not a failure of the theory, in my opinion. An adequate account of the term 'know' should make the temptations of skepticism comprehensible, which my theory does. But it should also put skepticism in a proper perspective, which Cartesianism fails to do.

In any event, I put forward my account of perceptual knowledge

not primarily as an antidote to skepticism, but as a more accurate rendering of what the term 'know' actually means. In this respect it is instructive to test my theory and its rivals against certain metaphorical or analogical uses of 'know'. A correct definition should be able to explain extended and figurative uses as well as literal uses, for it should explain how speakers arrive at the extended uses from the central ones. With this in mind, consider how tempting it is to say of an electric-eye door that it "knows" you are coming (at least that *something* is coming), or "sees" you coming. The attractiveness of the metaphor is easily explained on my theory: the door has a reliable mechanism for discriminating between something being before it and nothing being there. It has a "way of telling" whether or not something is there: this "way of telling" consists in a mechanism by which objects in certain DOE relations to it have differential effects on its internal state. By contrast, note how artificial it would be to apply more traditional analyses of 'know' to the electric-eye door, or to other mechanical detecting devices. How odd it would be to say that the door has "good reasons," "adequate evidence," or "complete justification" for thinking something is there; or that it has "the right to be sure" something is there. The oddity of these locutions indicates how far from the mark are the analyses of 'know' from which they derive.

The trouble with many philosophical treatments of knowledge is that they are inspired by Cartesian-like conceptions of justification or vindication. There is a consequent tendency to overintellectualize or overrationalize the notion of knowledge. In the spirit of naturalistic epistemology,¹⁷ I am trying to fashion an account of knowing that focuses on more primitive and pervasive aspects of cognitive life, in connection with which, I believe, the term 'know' gets its application. A fundamental facet of animate life, both human and infra-human, is telling things apart, distinguishing predator from prey, for example, or a protective habitat from a threatening one. The concept of knowledge has its roots in this kind of cognitive activity.

ALVIN I. GOLDMAN

The University of Michigan

¹⁷ Cf. W. V. Quine, "Epistemology Naturalized," in *Ontological Relativity, and Other Essays* (New York: Columbia, 1969).