10 Gettier was replying to Harman at a symposium of the American Philosophical Association, Eastern Division, 1970.

11 Chisholm, *op. cit.*, 105.

12 Skyrms, Brian, "The Explication of 'X knows that *p*'," *Journal of Philosophy*, 64 12 (1967), 373–389.

Knowledge

ROBERT NOZICK

CONDITIONS FOR KNOWLEDGE

Our task is to formulate further conditions to go alongside

(1) p is true
(2) S believes that p.

We would like each condition to be necessary for knowledge, so any case that fails to satisfy it will not be an instance of knowledge. Furthermore, we would like the conditions to be jointly sufficient for knowledge, so any case that satisfies all of them will be an instance of knowledge. We first shall formulate conditions that seem to handle ordinary cases correctly, classifying as knowledge cases which are knowledge, and as nonknowledge cases which are not; then we shall check to see how these conditions handle some difficult cases discussed in the literature.[1]

The casual condition on knowledge, previously mentioned, provides an inhospitable environment for mathematical and ethical knowledge; also there are well-known difficulties in specifying the type of casual connection. If someone floating in a tank oblivious to everything around him is given (by direct electrical and chemical stimulation of the brain) the belief that he is floating in a tank with his brain being stimulated, then even though that fact is part of the cause of his belief, still he does not know that it is true.

Let us consider a different third condition:

(3) If p weren't true, S wouldn't believe that p.

Throughout this work, let us write the subjunctive 'if-then' by an arrow, and the negation of a sentence by prefacing "not-" to it. The above condition thus is rewritten as:

(3) not-$p \rightarrow$ not-(S believes that p).

This subjunctive condition is not unrelated to the causal condition. Often when the fact that p (partially) causes someone to believe that p, the fact also will be causally necessary for his having the belief—without the cause, the effect would not occur. In that case, the subjunctive condition 3 also will be satisfied. Yet this condition is not equivalent to the causal condition. For the causal condition will be satisfied in cases of causal overdetermination, where either two sufficient causes of the effect actually operate, or a back-up cause (of the same effect) would operate if the first one didn't; whereas the subjunctive condition need not hold for these cases.[2] When the two conditions do agree, causality indicates knowledge because it acts in a manner that makes the subjunctive 3 true.

Reprinted by permission of the author and the publishers from *Philosophical Explanations* by Robert Nozick, Cambridge, Mass.: Harvard University Press

The subjunctive condition 3 serves to exclude cases of the sort first described by Edward Gettier, such as the following. Two other people are in my office and I am justified on the basis of much evidence in believing the first owns a Ford car; though he (now) does not, the second person (a stranger to me) owns one. I believe truly and justifiably that someone (or other) in my office owns a Ford car, but I do not know someone does. Concluded Gettier, knowledge is not simply justified true belief.

The following subjunctive, which specifies condition 3 for this Gettier case, is not satisfied: if no one in my office owned a Ford car, I wouldn't believe that someone did. The situation that would obtain if no one in my office owned a Ford is one where the stranger does not (or where he is not in the office); and in that situation I still would believe, as before, that someone in my office does own a Ford, namely, the first person. So the subjunctive condition 3 excludes this Gettier case as a case of knowledge.

The subjunctive condition is powerful and intuitive, not so easy to satisfy, yet not so powerful as to rule out everything as an instance of knowledge. A subjunctive conditional "if p were true, q would be true", $p \rightarrow q$, does not say that p entails q or that it is logically impossible that p yet not-q. It says that in the situation that would obtain if p were true, q also would be true. This point is brought out especially clearly in recent 'possible-worlds' accounts of subjunctives: the subjunctive is true when (roughly) in all those worlds in which p holds true that are closest to the actual world, q also is true. (Examine those worlds in which p holds true closest to the actual world, and see if q holds true in all these.) Whether or not q is true in p worlds that are still farther away from the actual world is irrelevant to the truth of the subjunctive. I do not mean to endorse any particular possible-worlds account of subjunctives, nor am I committed to this type of account.[3] I sometimes shall use it, though, when it illustrates points in an especially clear way.[4]

The subjunctive condition 3 also handles nicely cases that cause difficulties for the view that you know that p when you can rule out the relevant alternatives to p in the context. For, as Gail Stine writes, "what makes an alternative relevant in one context and not another? . . . if on the basis of visual appearances obtained under optimum conditions while driving through the countryside Henry identifies an object as a barn, normally we say that Henry knows that it is a barn. Let us suppose, however, that unknown to Henry, the region is full of expertly made papier-mâché facsimiles of barns. In that case, we would not say that Henry knows that the object is a barn, unless he has evidence against it being a papier-mâché facsimile, which is now a relevant alternative. So much is clear, but what if no such facsimiles exist in Henry's surroundings, although they once did? Are either of these circumstances sufficient to make the hypothesis (that it's a papier-mâché object) relevant? Probably not, but the situation is not so clear."[5] Let p be the statement that the object in the field is a (real) barn, and q the one that the object in the field is a papier-mâché barn. When papier-mâché barns are scattered through the area, if p were false, q would be true or might be. Since in this case (we are supposing) the person still would believe p, the subjunctive

(3) not-$p \rightarrow$ not-(S believes that p)

is not satisfied, and so he doesn't know that p. However, when papier-mâché barns are or were scattered around another country, even if p were false q wouldn't be true, and so (for all we have been told) the person may well know that p. A hypothesis q contrary to p clearly is relevant when if p weren't true, q would be true; when not-$p \rightarrow q$. It clearly is irrelevant when if p weren't true, q also would not be true; when not-$p \rightarrow$ not-q. The remaining possibility is that neither of these opposed subjunctives holds; q might (or might not) be true if p weren't true. In this case, q also will be relevant, according to an account of knowledge incorporating condition 3 and treating subjunctives along the lines sketched above. Thus, condition 3 handles cases that befuddle the "relevant alternatives" account; though that account can adopt the above subjunctive criterion for when an alternative is relevant, it then becomes merely an alternate and longer way of stating condition 3.[6]

Despite the power and intuitive force of the condition that if p weren't true the person would not believe it, this condition does not (in conjunction with the first two conditions) rule out every problem case. There remains, for example, the case of the person in the tank who is brought to believe, by direct electrical and chemical stimulation of his brain, that he is in the tank and is being brought to believe things in this way; he does not know this is true. However, the subjunctive condition is satisfied: if he weren't floating in the tank, he wouldn't believe he was.

The person in the tank does not know he is there, because his belief is not sensitive to the truth. Although it is caused by the fact that is its content, it is not sensitive to that fact. The operators of the tank could have produced any belief, including the false belief that he wasn't in the tank; if they had, he would have believed that. Perfect sensitivity would involve beliefs and facts varying together. We already have one portion of that variation, subjunctively at least: if p were false he wouldn't believe it. This sensitivity as specified by a subjunctive does not have the belief vary with the truth or falsity of p in all possible situations, merely in the ones that would or might obtain if p were false.

The subjunctive condition

(3) not-$p \rightarrow$ not-(S believes that p)

tells us only half the story about how his belief is sensitive to the truth-value of p. It tells us how his belief state is sensitive to p's falsity, but not how it is sensitive to p's truth; it tells us what his belief state would be if p were false, but not what it would be if p were true.

To be sure, conditions 1 and 2 tell us that p is true and he does believe it, but it does not follow that his believing p is sensitive to p's being true. This additional sensitivity is given to us by a further subjunctive: if p were true, he would believe it.

(4) $p \rightarrow$ S believes that p.

Not only is p true and S believes it, but if it were true he would believe it. Compare: not only was the photon emitted and did it go to the left, but (it was then true that): if it were emitted it would go to the left. The truth of antecedent and consequent is not alone sufficient for the truth of a subjunctive; 4 says more than 1 and 2.[7] Thus, we presuppose some (or another) suitable account of subjunctives. According to the suggestion tentatively made above, 4 holds true if not only does he actually truly believe p, but in the "close" worlds where p is true, he also believes it. He believes that p for some distance out in the p neighborhood of the actual world; similarly, condition 3 speaks not of the whole not-p neighborhood of the actual world, but only of the first portion of it. (If, as is likely, these explanations do not help, please use your own intuitive understanding of the subjunctives 3 and 4.)

The person in the tank does not satisfy the subjunctive condition 4. Imagine as actual a world in which he is in the tank and is stimulated to believe he is, and consider what subjunctives are true in that world. It is not true of him there that if he were in the tank he would believe it; for in the close world (or situation) to his own where he is in the tank but they don't give him the belief that he is (much less instill the belief that he isn't) he doesn't believe he is in the tank. Of the person actually in the tank and believing it, it is not true to make the further statement that if he were in the tank he would believe it—so he does not know he is in the tank.[8]

The subjunctive condition 4 also handles a case presented by Gilbert Harman.[9] The dictator of a country is killed; in their first edition, newspapers print the story, but later all the country's newspapers and other media deny the story, falsely. Everyone who encounters the denial believes it (or does not know what to believe and so suspends judgment). Only one person in the country fails to hear any denial and he continues to believe the truth. He satisfies conditions 1 through 3 (and the causal condition about belief) yet we are reluctant to say he knows the truth. The reason is that if he had heard the denials, he too would have believed them, just like everyone else. His belief is not sensitively tuned to the truth, he doesn't satisfy the condition that if it were true he would believe it. Condition 4 is not satisfied.[10]

There is a pleasing symmetry about how this account of knowledge relates conditions 3 and 4, and connects them to the first two conditions. The account has the following form.

(1)
(2)
(3) not-1 $\rightarrow$ not-2
(4) 1 $\rightarrow$ 2

I am not inclined, however, to make too much of this symmetry, for I found also that with other conditions experimented with as a possible fourth condition there was some way to construe the resulting third and fourth conditions as symmetrical answers to some symmetrical looking questions, so that they appeared to arise in parallel fashion from similar questions about the components of true belief.

Symmetry, it seems, is a feature of a mode of presentation, not of the contents presented. A uniform transformation of symmetrical statements can leave the results nonsymmetrical. But if symmetry attaches to mode of presentation, how can it possibly be a deep feature of, for instance, laws of nature that they exhibit symmetry? (One of my favorite examples of symmetry is due to Groucho Marx. On his radio program he spoofed a commercial, and ended, "And if you are not completely satisfied, return the unused portion of our product and we will return the unused portion of your money.") Still, to present our subject symmetrically makes the connection of knowledge to true belief especially perspicuous. It seems to me that a symmetrical formulation is a sign of our understanding, rather than a mark of truth. If we cannot understand an asymmetry as arising from an underlying symmetry through the operation of a particular factor, we will not understand why that asymmetry exists in that direction. (But do we also need to understand why the underlying asymmetrical factor holds instead of its opposite?)

A person knows that *p* when he not only does truly believe it, but also would truly believe it and wouldn't falsely believe it. He not only actually has a true belief, he subjunctively has one. It is true that *p* and he believes it; if it weren't true he wouldn't believe it, and if it were true he would believe it. To know that *p* is to be someone who would believe it if it were true, and who wouldn't believe it if it were false.

It will be useful to have a term for this situation when a person's belief is thus subjunctively connected to the fact. Let us say of a person who believes that *p*, which is true, that when 3 and 4 hold, his belief *tracks* the truth that *p*. To know is to have a belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it.

One refinement is needed in condition 4. It may be possible for someone to have contradictory beliefs, to believe *p* and also believe not-*p*. We do not mean such a person to easily satisfy 4, and in any case we want his belief-state, sensitive to the truth of *p*, to focus upon *p*. So let us rewrite our fourth condition as:

(4) $p \rightarrow$ S believes that p and not-(S believes that not-p).[11]

As you might have expected, this account of knowledge as tracking requires some refinements and epicycles. Readers who find themselves (or me) bogged down in these refinements should move on directly to this essay's second part, on skepticism, where the pace picks up. [See Chapter 5.]

WAYS AND METHODS

The fourth condition says that if *p* were true the person would believe it. Suppose the person only happened to see a certain event or simply chanced on a book describing it. He knows it occurred. Yet if he did not happen to glance that way or encounter the book, he would not believe it, even though it occurred. As written, the fourth condition would exclude this case as one where he actually knows the event occurred. It also would exclude the following case. Suppose some person who truly believes that *p* would or might arrive at a belief about it in some other close situation where it holds true, in a way or by a method different from the one he (actually) used in arriving at his belief that *p*, and so thereby comes to believe that not-*p*. In that (close) situation, he would believe not-*p* even though *p* still holds true. Yet, all this does not show he actually doesn't know that *p*, for actually he has not used this alternative method in arriving at his belief. Surely he can know that *p*, even though condition 4, as written, is not satisfied.

Similarly, suppose he believes that *p* by one method or way of arriving at belief, yet if *p* were false he wouldn't use this method but would use another one instead, whose application would lead him mistakenly to believe *p* (even though it is false). This person does not satisfy condition 3 as written; it is not true of him that if *p* were false he wouldn't believe it. Still, the fact that he would use another method of arriving at belief if *p* were false does not show he didn't know that *p* when he used this method. A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. Clearly, we must restate our conditions to take explicit account of the ways and methods of arriving at belief.

Let us define a technical locution, S knows, via method (or way of believing) M, that *p*:

(1) *p* is true.
(2) S believes, via method or way of coming to believe M, that *p*.
(3) If *p* weren't true and S were to use M to arrive at a belief whether (or not) *p*, then S wouldn't believe, via M, that *p*.
(4) If *p* were true and S were to use M to arrive at a belief whether (or not) *p*, then S would believe, via M, that *p*.

We need to relate this technical locution to our ordinary notion of knowledge. If only one method M is actually or subjunctively relevant to S's belief that *p*, then, simply, S knows that *p* (according to our ordinary notion) if and only if that method M is such that S knows that *p* via M.

Some situations involve multiple methods, however.

First Situation: S's belief that *p* is overdetermined; it was introduced (or reinforced) by two methods, each of which in isolation would have been sufficient to produce in S the belief that *p*. S's belief that *p* via one of these methods satisfies conditions 1–4. However, S's belief that *p* via the second method does not satisfy conditions 1–4, and in particular violates condition 3.

A case of this sort is discussed by Armstrong.[12] A father believes his son innocent of committing a particular crime, both because of faith in his son and (now) because he has seen presented in the courtroom a conclusive demonstration of his son's innocence. His belief via the method of courtroom demonstration satisfies 1–4, let us suppose, but his faith-based belief does not. If his son were guilty, he would still believe him innocent, on the basis of faith in his son. Thus, his belief that *p* (that his son is innocent) via faith in his son violates condition 3. Looking at his belief alone, without mention of method, his belief that *p* violates the third condition (namely, if *p* were false S wouldn't believe that *p*), which made no mention of method.

> *Second Situation:* S's belief that p via one method satisfies conditions 1–4. However, if p were false, S would not use that method in arriving at a belief about the truth value of p. Instead, he would use another method, thereby deciding, despite p's falsity, that p was true. S's actual belief that p is in no way based on the use of this second method, but if p were false he would believe p via the second method. (However, if p were false and S were to decide about its truth value by using the first method, then S would not believe that p. To be sure, if p were false S wouldn't decide about it by using that first method.) The truth value of p affects which method S uses to decide whether p.

Our earlier example of the grandmother is of this sort. Consider one further example, suggested to me by Avishai Margalit. S believes a certain building is a theater and concert hall. He has attended plays and concerts there (first method). However, if the building were not a theater, it would have housed a nuclear reactor that would so have altered the air around it (let us suppose) that everyone upon approaching the theater would have become lethargic and nauseous, and given up the attempt to buy a ticket. The government cover story would have been that the building was a theater, a cover story they knew would be safe since no unmedicated person could approach through the nausea field to discover any differently. Everyone, let us suppose, would have believed the cover story; they would have believed that the building they saw (but only from some distance) was a theater.

S believes the building is a theater because he has attended plays and concerts inside. He does not believe it is a theater via the second method of reading the government's cover story plus planted spurious theater and concert reviews. There are no such things. However, if it weren't a theater, it would be a nuclear reactor, there would be such cover stories, and S would believe still (this time falsely and via the second method) that the building was a theater. Nonetheless, S, who actually has attended performances there, knows that it is a theater.

To hold that a person knows that p if there exists at least one method M, satisfying conditions 1–4, via which he believes that p, would classify the father as knowing his son is innocent, a consequence too charitable to the father. Whereas it seems too stringent to require that all methods satisfy conditions 1–4, including those methods that were not actually used but would be under some other circumstances; the grandmother knows her grandson is well, and the person who has attended the concerts and plays knows the building is a theater. It is more reasonable to hold he knows that p if all the methods via which he actually believes that p satisfy conditions 1–4. Yet suppose our theatergoer also believes it is a theater partly because government officials, before they decided on which use they would put the building to, announced they were building a theater. Still, the theatergoer knows the building is a theater. Not all methods actually used need satisfy conditions 1–4, but we already have seen how the weak position that merely one such method is enough mishandles the case of the father.

We are helped to thread our way through these difficulties when we notice this father does not merely believe his son is innocent via the route of faith in his son; this defective route, not satisfying 1–4, also outweighs for him the method of courtroom demonstration. Even if courtroom demonstration (had it operated alone) would lead to the belief that his son is guilty, that not-p, still he would believe his son innocent, via faith in his son. Although it is the method of courtroom demonstration that gives him knowledge that p if anything does, for the father this method is outweighed by faith.[13] As a first try at delineating outweighing, we might say that method M is outweighed by others if when M would have the person believe p, the person believes not-p if the other methods would lead to the belief that not-p, or when M would have the person believe not-p, the person believes p if the other methods would lead to the belief that p.

This leads us to put forth the following position: S knows that p if there is some method via which S believes that p which satisfies conditions 1–4, and that method is not outweighed by any other method(s), via which S actually believes that p, that fail to satisfy conditions 3 and 4. According to this position, in some cases a person has knowledge even when he also actually believes via a method M_1 that does not satisfy 1–4, provided it is outweighed by one that does; namely, in the overdetermination case, and in the case when M_1 alone would suffice to fix belief but only in the absence of a verdict from the M he also uses which does satisfy 1–4.

S knows that p if and only if there is a method M such that (a) he knows that p via M, his belief via M that p satisfies conditions 1–4, and (b) all other methods M_1 via which he believes that p that do not satisfy conditions 1–4 are outweighed by M.[14]

We have stated our outweighing requirement only roughly; now we must turn to refinements. According to our rough statement, in the overdetermination case, method M_1, which satisfies 3 and 4 and which is what gives knowledge if anything does, wins out over the other method M_2 in all cases. The actual situation (Case I) is where M_1 recommends believing p as does M_2, and the person believes p. In this case we have made our answer to the question whether he knows that p depend on what happens or would happen in the two other cases where the methods recommend different beliefs. (See Table.) The first rough statement held that the person knows in Case I only if he would believe p in Case II and not-p in Case III. While this is sufficient for knowledge in Case I, it seems too stringent to be necessary for such knowledge.

An alternative and more adequate view would hold constant what the other method recommends, and ask whether the belief varies with the recommendation of M_1. Since M_2 actually recommends p (Case I), we need look only at Case III and ask: when M_2 continues to recommend p and M_1 recommends not-p, would the person believe not-p? Despite his faith, would the father believe his son guilty if the courtroom procedure proved guilt? That is the relevant question—not what he would believe if the courtroom showed innocence while (somehow) his method of faith led to a conclusion of guilty.

Consider how this works out in another simple case. I see a friend today; he is now alive. However, if he were not alive, I wouldn't have seen him today or (let us suppose) heard of his death, and so still would believe he was alive. Yet condition 3 is satisfied; it includes reference to a method, and the method M_1 of seeing him satisfies 3 with respect to p equals he is alive at the time. But there also is another method M_2 via which I believe he is alive, namely having known he was alive yesterday and continuing to believe it. Case III asks what I would believe if I saw the friend dead (though I knew yesterday he was alive); our position holds I must believe him dead in this case if I am to know by seeing him that he is alive in Case I. However, we need not go so far as to consider what I would believe if I had "learned" yesterday that he was dead yet "saw" him alive today. Perhaps in that case I would wonder whether it really was he I was seeing. Even so, given the result in Case III, I know (in Case I) he is alive. Thus, we hold fixed the recommendation of the other method, and only ask whether then the belief varies with the recommendation of method M_1.[15]

Our test of looking at Case III cannot apply if M_1 is a one-sided method, incapable of recommending belief in not-p; it either recommends belief in p or yields no recommendation. (Perhaps M_1 detects one of a

	M_1 Recommends	*M_2 Recommends*	*Does the Person Believe* p *or Believe not-*p?
Case I	believe p	believe p	believes p
Case II	believe p	believe not-p	?
Case III	believe not-p	believe p	?

number of sufficient conditions for p; not detecting this, M_1 remains silent as to the truth of p.) What are we to say about his knowing if a person's belief is overdetermined or jointly determined by a one-sided method M_1 plus another method M_2 which fails to satisfy condition 3? Should we now look at Case II, where M_1 recommends belief in p and M_2 recommends belief in not-p, and say that believing p in this case is sufficient to show that M_1 outweighs M_2? That does not seem unreasonable, but we had better be careful to stipulate that this Case II situation is a sufficient condition for M_1's outweighing M_2 only when the Case III situation is impossible, for otherwise we face the possibility of divergent results. (For example, he believes p in Case II and in Case III, yet believes not-p when both methods recommend not-p; here the result in Case II indicates M_1 outweighs M_2 while the result in Case III indicates M_2 outweighs M_1.) It is Case III that should predominate.

One final remark about method. Suppose a method is good for some types of statements but not others; it satisfies 3 and 4 for the first type but not for the second. However, S believes the method is good for all types of statements and applies it indiscriminately. When he applies it to a statement of the first type which he thereby comes to believe, does he know that it is true? He does, if he satisfies conditions 3 and 4. Hesitation to grant him knowledge stems, I think, from the fact that if p were false and were of the second type, he might well still believe it. Whether or not this undercuts condition 3 for knowledge depends upon the disparity of the two types; the greater the gulf between the types, the more willing we are to say he knows a statement of the type where M works.

In explaining the nature of knowledge by reference to a method or way of believing, we leave large questions open about how to individuate methods, count them, identify which method is at work, and so on. I do not want to underestimate these difficulties, but neither do I want to pursue them here.[16] Still, some clarifying remarks are needed.

A person can use a method (in my sense) without proceeding methodically, and without knowledge or awareness of what method he is using. Usually, a method will have a final upshot in experience on which the belief is based, such as visual experience, and then (a) no method without this upshot is the same method, and (b) any method experientially the same, the same "from the inside", will count as the same method. Basing our beliefs on experiences, you and I and the person floating in the tank are using, for these purposes, the same method.

Some methods are supervenient on others, for example, "believing what seems to be true to you" or "believing what seems true given the weighting of all other methods". The account of outweighing is not to apply to such supervenient methods, otherwise there always will be such a one that outweighs all the others. There are various gerrymandered (Goodmanesque) methods that would yield the same resulting belief in the actual situation; which method a person actually is using will depend on which general disposition to acquire beliefs (extending to other situations) he actually is exercising.[17]

Although sometimes it will be necessary to be explicit about the methods via which someone believes something, often it will cause no confusion to leave out all mention of method. Furthermore, some statements play a central role in our continuing activities, or in our picture of the world or framework wherein we check other statements, for example, "I have two hands", "the world has existed for many years already"; it is misleading to think of our coming to believe them via some delimited method or methods.[18] So nested are these statements in our other beliefs and activities, and so do they nest them, that our belief or acceptance of them is (for almost all purposes) best represented apart from any particular methods. In considering our knowledge of them we may revert to the earlier simpler subjunctives

(3) not-p → not-(S believes that p)

(4) p → S believes that p.

The very centrality of the specific *p* means that 4 will be satisfied without reference to a specific method or way of believing. In contrast, I know there is a pair of scissors on my desk (in front of me) now; but it is not accurate simply to say that if there were a pair of scissors there, I would believe there was. For what if I weren't looking, or hadn't looked, or were elsewhere now? Reference to the method via which I believe there are scissors on the desk is needed to exclude these possibilities. With the most central statements, however, there is no similar "what if"; their centrality ensures they will not escape notice.

NOTES

1. Despite some demurrals in the literature, there is general agreement that conditions 1 and 2 are necessary for knowledge. (For some recent discussions, see D. M. Armstrong, *Belief, Truth and Knowledge*, Cambridge University Press, 1973, ch. 10; Keith Lehrer, *Knowledge*, Oxford University Press, 1974, chs. 2, 3.) I shall take for granted that this is so, without wishing to place very much weight on its being belief that is the precise cognitive attitude (as opposed to thinking it so, accepting the statement, and so on) or on the need to introduce truth as opposed to formulating the first condition simply as: *p*.

I should note that our procedure here does not stem from thinking that every illuminating discussion of an important philosophical notion must present (individually) necessary and (jointly) sufficient conditions.

2. Below, we discuss further the case where though the fact that *p* causes the person's belief that *p*, he would believe it anyway, even if it were not true. I should note here that I assume bivalence throughout this chapter, and consider only statements that are true if and only if their negations are false.

3. See Robert Stalnaker, "A Theory of Conditionals", in N. Rescher, ed., *Studies in Logical Theory* (Basil Blackwell, Oxford, 1968); David Lewis, *Counterfactuals* (Harvard University Press, Cambridge, 1973); and Jonathan Bennett's critical review of Lewis, "Counterfactuals and Possible Worlds", *Canadian Journal of Philosophy*, Vol. IV, no. 2, Dec. 1974, pp. 381–402.

Our purposes require, for the most part, no more than an intuitive understanding of subjunctives. However, it is most convenient to examine here some further issues, which will be used once or twice later. Lewis' account has the consequence that $p \rightarrow q$ whenever *p* and *q* are both true; for the possible world where *p* is true that is closest to the actual world is the actual world itself, and in that world *q* is true. We might try to remedy this by saying that when *p* is true, $p \rightarrow q$ is true if and only if *q* is true in all *p* worlds closer (by the metric) to the actual world than is any not-*p* world. When *p* is false, the usual accounts hold that $p \rightarrow q$ is true when *q* holds merely in the closest *p* worlds to the actual world. This is too weak, but how far out must one go among the *p* worlds? A suggestion parallel to the previous one is: out until one reaches another not-*p* world (still further out). So if *q* holds in the closest *p* world w_1, but not in the *p* world w_2, even though no not-*p* world lies between w_1 and w_2, then (under the suggestion we are considering) the subjunctive is false. A unified account can be offered for subjunctives, whatever the truth value of their antecedents. The *p* neighborhood of the actual world A is the closest *p* band to it; that is, w is in the *p* neighborhood of the actual world if and only if *p* is true in w and there are no worlds w^{p} and w^{p} such that not-*p* is true in w^{p} and *p* is true in w^{p}, and w^{p} is closer to A than w is to A, and w^{p} is at least as close to A as w^{p} is to A. A subjunctive $p \rightarrow q$ is true if and only if *q* is true throughout the *p* neighborhood of the actual world.

If it is truly a random matter which slit a photon goes through, then its going through (say) the right slit does not establish the subjunctive: if a photon were fired at that time from that source it would go through the right-hand slit. For when *p* equals A photon is fired at that time from that source, and *q* equals the photon goes through the right-hand slit, *q* is not true everywhere in the *p* neighborhood of the actual world.

This view of subjunctives within a possible-worlds framework is inadequate if there is no discrete *p* band of the actual world, as when for each positive distance from the actual world A, there are both *p* worlds and not-*p* worlds so distant. Even if this last is not generally so, many *p* worlds that interest us may have their distances from A matched by not-*p* worlds. Therefore, let us redefine the relevant *p* band as the closest spread of *p* worlds such that there is no not-*p* world intermediate in distance from A to two *p* worlds in the spread unless there is also another *p* world in the spread the very same distance from A. By definition, it is only *p* worlds in the *p* band, but some not-*p* worlds may be equidistant from A.

Though this emendation allows us to speak of the closest spread of *p* worlds, it no longer is so clear which worlds in this *p* band subjunctives (are to) encompass. We have said it is not sufficient for the truth of $p \rightarrow q$ that *q* hold in that one world in the *p* band closest to the actual world. Is it necessary, as our first suggestion has it, that *q* hold in all the *p* worlds in the closest *p* band to the actual world? Going up until the first "pure" stretch of not-*p* worlds is no longer as natural a line to draw as when we imagined "pure" *p* neighborhoods. Since there already are some not-*p* worlds the same distance from A as some members of the *p* band, what is the special significance of the first unsullied not-*p* stretch? There seems to be no natural line, though, coming before this stretch yet past the first *p* world. Perhaps nothing stronger can be said than this: $p \rightarrow q$ when *q* holds for some distance out in the closest *p* band to the actual world, that is, when all the worlds in this first part of that closest p band are q. The distance need not be fixed as the same for all subjunctives, although various general formulas might be imagined, for example, that the distance is a fixed percentage of the width of the p band.

I put forth this semantics for subjunctives in a possible-worlds framework with some diffidence, having little inclination to pursue the details. Let me emphasize, though, that this semantics does not presuppose any realist view that all possible worlds obtain. (Such a view was discussed in the previous chapter.) I would hope that into this chapter's subjunctively formulated theoretical structure can be plugged (without too many modifications) whatever theory of subjunctives turns out to be adequate, so that the theory of knowledge we formulate is not sensitive to variations in the analysis of subjunctives. In addition to Lewis and Stalnaker cited above, see Ernest W. Adams, *The Logic of Conditionals* (Reidel, Dodrecht, 1975); John Pollock, *Subjunctive Reasoning* (Reidel, Dodrecht, 1976); J. H. Sobel, "Probability, Chance and Choice" (unpublished book manuscript); and a forthcoming book by Yigal Kvart.

4. If the possible-worlds formalism is used to represent counterfactuals and subjunctives, the relevant worlds are not those p worlds that are closest or most similar to the actual world, unless the measure of closeness or similarity is: what would obtain if p were true. Clearly, this cannot be used to explain when subjunctives hold true, but it can be used to represent them. Compare utility theory which represents preferences but does not explain them. Still, it is not a trivial fact that preferences are so structured that they can be represented by a real-valued function, unique up to a positive linear transformation, even though the representation (by itself) does not explain these preferences. Similarly, it would be of interest to know what properties hold of distance metrics which serve to represent subjunctives, and to know how subjunctives must be structured and interrelated so that they can be given a possible-worlds representation. (With the same one space serving for all subjunctives?)

One further word on this point. Imagine a library where a cataloguer assigns call numbers based on facts of sort F. Someone, perhaps the cataloguer, then places each book on the shelf by looking at its call number, and inserting it between the two books whose call numbers are most nearly adjacent to its own. The call number is derivative from facts of type F, yet it plays some explanatory role, not merely a representational one. "Why is this book located precisely there? Because of its number." Imagine next another library where the person who places books on the shelves directly considers facts of type F, using them to order the books and to interweave new ones. Someone else might notice that this ordering can be represented by an assignment of numbers, numbers from which other information can be derived as well, for example, the first letter of the last name of the principal author. But such an assigned number is no explanation of why a book in this library is located between two others (or why its author's last name begins with a certain letter). I have assumed that utility numbers stand to preferences, and closeness or similarity measures stand to subjunctives, as the call numbers do to the books, and to the facts of type F they exhibit, in the second library.

5. G. C. Stine, "Skepticism, Relevant Alternatives and Deductive Closure", *Philosophical Studies*, Vol. 29, 1976, p. 252, who attributes the example to Carl Ginet.

6. This last remark is a bit too brisk, for that account might use a subjunctive criterion for when an alternative q to p is relevant (namely, when if p were not to hold, q would or might), and utilize some further notion of what it is to rule out relevant alternatives (for example, have evidence against them), so that it did not turn out to be equivalent to the account we offer.

7. More accurately, since the truth of antecedent and consequent is not necessary for the truth of the subjunctive either, 4 says something different from 1 and 2.

8. I experimented with some other conditions which adequately handled this as well as some other problem cases, but they succumbed to further difficulties. Though much can be learned from applying those conditions, presenting all the details would engage only the most masochistic readers. So I simply will list them, each at one time a candidate to stand alone in place of condition 4.

(a) S believes that not-p → not-p.
(b) S believes that not-p → not-p or it is through some other method that S believes not-p. (Methods are discussed in the next section.)
(c) (S believes p or S believes not-p) → not-(S believes p, and not-p holds) and not-(S believes not-p, and p holds).
(d) not-(S believes that p) → not-(p and S believes that not-p).
(e) not-(p and S believes that p) → not-(not-p and S believes that p or p and S believes that not-p).

9. Gilbert Harman, *Thought* (Princeton University Press, Princeton, 1973), ch. 9, pp. 142–154.

10. What if the situation or world where he too hears the later false denials is not so close, so easily occurring? Should we say that everything that prevents his hearing the denial easily could have not happened, and does not in some close world?

11. This reformulation introduces an apparent asymmetry between the consequents of conditions 3 and 4.

Since we have rewritten 4 as

p → S believes that p and not-(S believes that not-p),

why is 3 not similarly rewritten as

not-p → not-(S believes that p) and S believes that not-p?

It is knowledge that p we are analyzing, rather than knowledge that not-p. Knowledge that p involves a stronger relation to p than to not-p. Thus, we did not first write the third condition for knowledge of p as:

not-$p \rightarrow$ S believes that not-p; also the following is not true: S knows that $p \rightarrow$ (not-$p \rightarrow$ S knows that not-p).

Imagine that someone S knows whether or not p, but it is not yet clear to us which he knows, whether he knows that p or knows that not-p. Still, merely given that S knows that——, we can say:

not-$p \rightarrow$ not-(S believes that p)
$p \rightarrow$ not-(S believes that not-p).

Now when the blank is filled in, either with p or with not-p, we have to add S's believing it to the consequent of the subjunctive that begins with it. That indicates which one he knows. Thus, when it is p that he knows, we have to add to the consequent of the second subjunctive (the subjunctive that begins with p): S believes that p. We thereby transform the second subjunctive into:

$p \rightarrow$ not-(S believes that not-p) and S believes that p.

Except for a rearrangement of which is written first in the consequent, this is condition 4. Knowledge that p especially tracks p, and this special focus on p (rather than not-p) gets expressed in the subjunctive, not merely in the second condition.

There is another apparent asymmetry in the antecedents of the two subjunctives 3 and 4, not due to the reformulation. When actually p is true and S believes that p, condition 4 looks some distance out in the p neighborhood of the actual world, while condition 3 looks some distance out in the not-p neighborhood, which itself is farther away from the actual world than the p neighborhood. Why not have both conditions look equally far, revising condition 3 to require merely that the closest world in which p is false yet S believes that p be some distance from the actual world. It then would parallel condition 4, which says that the closest world in which p yet p is not believed is some distance away from the actual world. Why should condition 3 look farther from the actual world than condition 4 does?

However, despite appearances, both conditions look at distance symmetrically. The asymmetry is caused by the fact that the actual world, being a p world, is not symmetrical between p and not-p. Condition 3 says that in the closest not-p world, not-(S believes that p), and that this 'not-(S believes that p)' goes out through the first part of the not-p neighborhood of the actual world. Condition 4 says that in the closest p world, S believes that p, and that this 'S believes that p' goes out through the first part of the p neighborhood of the actual world. Thus the two conditions are symmetrical; the different distances to which they extend stems not from an asymmetry in the conditions but from one in the actual world—it being (asymmetrically) p.

12. D. M. Armstrong, *Belief, Truth and Knowledge* (Cambridge University Press, 1973), p. 209; he attributes the case to Gregory O'Hair.

13. Some may hold the father is made more sure in his belief by courtroom proof; and hold that the father knows because his degree of assurance (though not his belief) varies subjunctively with the truth.

14. If there is no other such method M_1 via which S believes that p, the second clause is vacuously true.

Should we say that no other method used outweighs M, or that M outweighs all others? Delicate questions arise about situations where the methods tie, so that no subjunctive holds about one always winning over the other. It might seem that we should require that M outweigh (and not merely tie) the other methods; but certain ways of resolving the ties, such as not randomly deciding but keeping judgment suspended, might admit knowledge when a true belief is arrived at via a tracking method M which is not outweighed yet also doesn't (always) outweigh the others present. There is no special need to pursue the details here; the outweighing condition should be read here and below as a vague one, residing somewhere in the (closed) interval between "outweighs" and "not outweighed", but not yet precisely located. This vagueness stands independently of the refinements pursued in the text immediately below.

15. When a belief is overdetermined or jointly produced by three methods, where only the first satisfies conditions 3 and 4, the question becomes: what does the person believe when M_1 recommends believing not-p while the two others each recommend believing p? Notice also that in speaking of what would happen in Case III we are imposing a subjunctive condition; if there is no "would" about it, if in each instance of a Case III situation it is determined at random which method outweighs which, then that will not be sufficient for knowledge, even though sometimes M_1 wins out.

It is worrisome that in weakening our initial description of outweighing by looking to Case III but not to Case II, we seem to give more weight to condition 3 for tracking than to condition 4. So we should be ready to reconsider this weakening.

16. For example, in the case of the father who believes on faith that his son is innocent and sees the courtroom demonstration of innocence, does the father use two methods, faith and courtroom demonstration, the second of which does satisfy conditions 3–4 while the first (which outweighs it) does not satisfy 3–4; or does the father use only one method which doesn't satisfy 3–4, namely: believe about one's son whatever the method of faith tells one, and only if it yields no answer, believe the result of courtroom demonstration? With either mode of individuation, knowledge requires the negative existentially quantified statement (that there is no method . . .) somewhere, whether in specifying the method itself or in specifying that it is not outweighed.

17. One suspects there will be some gimmick whereby whenever p is truly believed a trivial method M can be specified which satisfies conditions 3 and 4. If so, then further conditions will have to be imposed upon M, in addition to the dispositional condition. Compare the difficulties encountered in the literature on specifying the relevant reference class in probabilistic inference and explanation; see Henry Kyburg, *Probability and the Logic of Rational Belief* (Wesleyan University Press, Middletown, 1961), ch. 9; C. G. Hempel, *Aspects of Scientific Explanation* (Free Press, New York, 1965), pp. 394–405; also his "Maximal Specificity

and Lawlikeness in Probabilistic Explanation", *Philosophy of Science,* Vol. 35, 1968, pp. 116–133.

18. See Ludwig Wittgenstein, *On Certainty* (Basil Blackwell, Oxford, 1969),§§ 83, 94, 102–110, 140–144, 151–152, 162–163, 166, 411, 419, 472–475.

SELECTED BIBLIOGRAPHY

AUDI, ROBERT, *Belief, Justification, and Knowledge.* Belmont; Calif.: Wadsworth Publishing Company, Inc., 1988, ch. 7.

DRETSKE, FRED, "Conclusive Reasons," *Australasian Journal of Philosophy* 49 (1971), 1–22.

———, "The Pragmatic Dimension of Knowledge," *Philosophical Studies* 40 (1981), 363–78.

FELDMAN, RICHARD, "An Alleged Defect in Gettier Counter-Examples," *Australasian Journal of Philosophy* 49 (1971), 1–22.

GOLDMAN, ALVIN, *Epistemology and Cognition.* Cambridge; Mass.: Harvard University Press, 1986, ch. 3.

HARMAN, GILBERT, "Knowledge, Reasons and Causes," *The Journal of Philosophy* 67 (1970), 841–55.

———, *Thought.* Princeton, N.J.: Princeton University Press, 1973.

KAPLAN, MARK, "It's Not What You Know That Counts," *The Journal of Philosophy* 82 (1985), 350–63.

LEHRER, KEITH, *Theory of Knowledge.* Boulder, Colo.: Westview Press, Inc., 1991, ch. 1.

LEHRER, KEITH, and THOMAS D. PAXSON, JR., "Knowledge: Undefeated Justified True Belief," *The Journal of Philosophy* 66 (1969), 225–37.

O'CONNOR, D. J., and BRIAN CARR. *Introduction to the Theory of Knowledge,* Minneapolis, Minn.: University of Minnesota Press, 1982, ch. 3.

PAPPAS, GEORGE S., and MARSHALL SWAIN, *Essays on Knowledge and Justification.* Ithaca, N.Y.: Cornell University Press, 1978.

POLLOCK, JOHN L., *Contemporary Theories of Knowledge.* Totowa, N.J.: Rowman & Littlefield, Publishers, Inc., 1986, pp. 180–93.

SHOPE, ROBERT K., *The Analysis of Knowing.* Princeton, N.J.: Princeton University Press, 1983.

SOSA, ERNEST, *Knowledge in Perspective; Selected Essays in Epistemology.* Cambridge: Cambridge University Press, 1991, sec. 1.

UNGER, PETER, "The Cone Model of Knowledge," *Philosophical Topics* 14 (1986), 125–78.

WILLIAMS, MICHAEL, "Inference, Justification, and the Analysis of Knowledge," *The Journal of Philosophy* 75 (1978), 249–63.

WITTGENSTEIN, LUDWIG, *The Blue and Brown Books.* New York: Harper & Row, Publishers, Inc., 1965, pp. 17–20.

———. In *Philosophical Investigations,* 3rd edition, ed. G. E. M. Anscombe. New York: Macmillan Publishing Co., Inc., 1958, secs. 65–78.