

five

## NOZICK ON KNOWLEDGE: FINDING THE RIGHT CONNECTION

Alan H. Goldman

When Gettier attacked the equation of justified true belief with knowledge by appealing to true beliefs inferred from justified but false premises,<sup>1</sup> three types of response emerged in the subsequent literature. First, epistemologists attempted to specify a connection between the truth of the proposition believed (or the fact to which the belief referred) and the belief itself, so as to prevent the belief from being only accidentally true.<sup>2</sup> A second alternative was to beef up the justification condition so as to rule out misleading evidence. This alternative in turn divided into two main tactics. The first sought to specify and rule out defeaters of justification, true propositions not known to the subject that would remove his warrant for believing.<sup>3</sup> The second sought to guarantee justification by requiring reliability in the source or method of acquiring the belief.<sup>4</sup> (Advocates of this type of analysis think of reliability as justification. Whether this fits the usual notion of justification will be considered briefly below.)

I

Nozick's analysis of knowledge represented one culmination (at the time) of the first type of response to Gettier examples.<sup>5</sup> He requires that a belief be counterfactually related to the fact to which it refers in the following sense. S's belief that *p* counts as knowledge only when:

- (1) if *p* were not true, then S would not believe it;
- (2) if *p* were true (in nonactual close worlds), then S would believe it.

When these conditions hold, S is said to track the fact that *p*. This relation need not obtain over all possible worlds, only over the closest worlds in which the antecedents are true.

Nozick amends these conditions to accommodate examples such as that of the grandmother who sees her grandson and therefore knows he is well, but would be deceived by other relatives as to the grandson's health were he sick. Here she knows that her grandson is well even though the first condition above is violated: she would not believe him to be sick if he were. Nozick's amendment is to require that the source or method of knowing be held constant in the relevant counterfactual situations: if the grandmother were to see her sick grandson, then she would not believe him to be well. As so amended, the analysis assures that a knower's belief be nonaccidentally connected with the fact that makes it true (in a certain sense of 'nonaccidental'). Thus it rules out the standard Gettier counterexamples to the more traditional analysis of knowledge, in which a believer is only accidentally correct in his belief, having inferred it from a false premise.

Both tracking conditions turn out, however, to be both too strong and too weak. We may begin with the charge that the first condition is too strong, that it rules out genuine cases of knowledge. Nozick admits this when it comes to knowing that I am using a particular method for generating belief by using that very method (216).<sup>6</sup> I can sense that I am sensing, but cannot hold the method constant while considering the situation in which I am not sensing. Hence I cannot apply his criterion to determine that I know I am sensing by sensing that I am.

Consider next the following case, illustrative of a more general difficulty. I am playing tennis with my twelve-year-old son when I hear from a nearby radio broadcast that an assassination has occurred at a distant location. I know from seeing my son (while I hear the broadcast) that he is not the assassin. Let us suppose that in that very distant possible world in which I witness my son assassinate the political leader in question, I am so driven to distraction that I do not come to believe him guilty. That (counterfactual) fact is irrelevant to my present knowledge in the real world. What seems to make it irrelevant is the fact that the possible world in question is so very distant from this one, in which my son is young, generally well-behaved, politically unmotivated, and playing tennis with me at the time of the murder. In the absence of any preconceived analysis of knowledge, we certainly would not question my claim to know my son's innocence in the situation described. Indeed this knowledge would be almost too obvious and trivial to assert in conceivable circumstances.

An example of similar structure was suggested to me by Risto Hilpinen. I am looking at a thermometer that is accurate within the range of 0 degrees to 100 degrees F. At all temperatures below 0 the

thermometer registers 0 degrees. By observing its reading of 70 degrees I come to have the belief that it is not - 50 degrees. I thereby know that it is not - 50 degrees, even though if it were, I would not believe it to be. Once again we see that a measuring instrument which links an observer to facts within its range can produce knowledge when it generates true beliefs from observations in that range. (Think of the visual system as such an instrument in the previous example.) If the closest world in which such a fact fails to obtain is distant, then the lack of tracking in that world appears to be irrelevant to the knowledge claim. (On the other hand, if the thermometer is registering close to its threshold of accuracy and is not always accurate around that threshold, then my beliefs based on it regarding values just outside its range will not qualify as knowledge. Here the possible world in which the link between fact and belief is broken is close.)

Thus Nozick's first tracking condition is too strong. In other cases it proves to be too weak, failing to disallow bogus claims to knowledge. Such is the case when the truth of a proposition is counterfactually linked to a belief, but only via a (causal or evidential) chain intuitively considered abnormal and defective. Consider a standard Gettier case intentionally set up to deceive its subject. I have been tricked into believing that someone in my office owns a Ford by Haveit's having given me evidence that Nogot owns one, when Haveit is the owner. If my belief were not true, then Haveit would not have led me to hold it. Yet I do not know that someone in the office owns a Ford, since the usual Gettier conditions obtain. I have inferred my belief from a false premise, although I would not have come to hold it if it had been false.

While the previous examples show that the first counterfactual condition in itself is too strong, this one shows that, when it links a believer to a fact via a chain of evidence, some stronger constraint on such chains is required. Nozick himself in another section of his chapter on knowledge proposes a constraint on inference that transmits knowledge. Such inference must satisfy the following counterfactual condition: if the conclusion were not true, then the premises would not be believed. Belief in the premises must track the truth of the conclusion. But this constraint too seems to be met in the intentional Gettier case, since Haveit would not have provided me with evidence for a false conclusion in the example. Hence we need a different constraint on such chains.

It might be thought that this example could be avoided by appealing to the second of Nozick's tracking conditions, since there might be a close world in which Haveit owns a Ford but does not present the misleading evidence to me. But, for one thing, considering such a world fails to hold constant the method by which I arrive at belief in the example. That method consists in inference from the evidence with which I am presented.

Problems about holding the method constant in fact lead us to charges against the second tracking condition similar to those raised against the first. Nozick himself wishes to rule out as knowledge the case of a brain in a vat that is programmed to have the belief that it is a brain in a vat (175-76). He appeals to his second tracking condition, arguing that there is a close world in which the belief is true but the programmer does not enter it into the brain. Here the brain does not itself use a method for obtaining the belief, and so there is no method to hold constant in the counterfactual condition. This claim enables Nozick to contrast this case with another in which I happen to see an object and gain knowledge of its location by looking in a certain direction, when there is a close world in which I do not look in that direction or see the object. The latter possibility is irrelevant, since my method that must be held constant includes looking in a particular direction.

One problem here, however, is that this contrast breaks down given a slight modification in the first example. Suppose that the brain's programmer gives it sensations of being surrounded by a vat. Then it will have a method of arriving at the belief that there is a vat around it, namely trusting in its visual experience. This makes the case analogous to that of my seeing an object; yet, as in the original brain-in-the-vat case, the brain lacks knowledge of its surroundings. (While it can have knowledge of general propositions, it cannot, I take it, acquire knowledge of its surroundings by this method.) Nozick's second tracking condition fails to disallow the knowledge claim, since, in the close worlds in which the method is held constant, in which the brain continues to believe on the basis of its present visual experience, it continues to believe that there is a vat around it.<sup>7</sup> (We may also imagine that the first tracking condition is satisfied in this example, in that the programmer would not have caused this belief if it were not true.)

Thus the second condition is too weak. Like the first, it is also too strong, as the following example shows. Imagine an old logician with failing mental capacities. At this point he is still capable of deriving the correct conclusion in a particular proof, but there is a close world in which he can no longer do so. Since he is still capable at this time of deriving the correct conclusion, he knows it when he does so. Yet the second tracking condition is violated, even when the method, inferring from the premises in question, is held constant. The moral is similar to that in the previous cases in which the first condition was found to be too strong: a method that links a subject to facts within its range need not always do so outside that range. In this case relating to the second tracking condition the connection is broken in a close world, and yet the subject still appears to have knowledge. It seems that an entirely different sort of connection is required here, one that need not connect belief to fact in all close worlds.

## II

That counterexamples to a given analysis of knowledge can be found is not sufficient reason in itself to give it up. We must allow tradeoffs between the simplicity of an account, its ability to guide the epistemologist on the normative level, and its ability to capture our convictions in exotic cases. It may turn out that the addition of conditions sufficient to capture all our intuitions (if indeed these are consistent) will render an analysis *ad hoc*, unmanageable, and of no use in the inquiry into the sources and extent of our knowledge and into methods for the improvement of those sources. Nozick's account certainly has the virtue of simplicity and elegance. It remains to be seen how well it fares with more traditional epistemological concerns.

Unfortunately, his discussion of skepticism, while unquestionably ingenious in tracing the implications of the analysis for the attempt to meet skeptical challenges, leaves us less than fully satisfied as well. According to his account we cannot know that either pervasive or more local skeptical alternatives to our beliefs do not in fact obtain. Consider the theses that our everyday perceptual beliefs systematically err because we are deceived by Descartes' demon, or that we are all those brains in vats being programmed to have the experiences we do. If these skeptical possibilities were real, we would not believe they were. Hence we do not track the fact that they are not real. Hence we do not know that we are not brains in vats or dupes of Descartes' demon.

While Nozick's analysis forces this sweeping concession to the skeptic, it also saves our everyday knowledge from skeptical refutation, saves at least the possibility of such knowledge, by implying a lack of logical closure, even in cases of obvious implications known to hold. Because subjunctive counterfactuals in general are not closed under implication, neither is tracking. According to this analysis I can know that *p*, know that *p* implies *q*, draw the inference, and not know that *q* (205-6). I can know various propositions to be true without knowing the skeptic's alternatives to be false, even though the truth of the former implies the falsity of the latter, because I track the former but not the latter. I can track the truth of everyday perceptual beliefs without tracking the denial of the skeptical theses because the skeptic's worlds are not the closest ones in which the everyday propositions I believe are false. If they were false, I would not believe them, although in the skeptic's worlds they are false and I do believe them. If my pipe were not on my desk before me now, I would not believe it to be, even though if I were now a brain in a vat, I would falsely believe there to be a pipe on a desk before me.

Despite the plausible ring of examples such as that just cited, other more local concessions to the skeptic are far less plausible. When tracking is required and closure denied, I cannot know, for example,

that my son is not a robot brilliantly constructed by aliens, although I can know that I do not have a brilliantly constructed robot son. Although these propositions are not equivalent, intuitively there seems to be no distinction in my ability to know them to be true. According to Nozick's account the distinction between them is epistemically crucial, since in the closest world in which my son is a robot, I do not believe he is, while in the closest world in which I have a robot son, I know that I have one (219-20). We should prefer an analysis of knowledge that allows us to know that we are not brains in vats, that our children are not robots (or assassins), and so on.

We perhaps require here an additional qualification and reminder about the place of intuitions in epistemological analysis. Intuitions about what and when we know are not sacrosanct. Given a realist view of the world and our relation to it, a view that Nozick shares, such initial epistemic evaluations are especially vulnerable, since knowledge requires a relation between facts and beliefs about which we can be mistaken. On the other hand, we attempt in analysis to capture our shared concept of knowledge. We must therefore begin from our settled convictions as data, adjusting in the end to preserve consistency and to preserve what is important in guiding inquiry. By so doing we will capture a real relation if our theory of knowledge is a good theory and if most of our intuitions about when we know are indeed correct. (Showing their correctness is a matter of answering the skeptic.)

We must not beg questions by beginning from those intuitions that match our initially favored analysis and then dismiss others inconsistent with it. An analogous procedure is legitimate only when it becomes otherwise clear that our convictions fail to cohere internally (or with any analysis derived partially from them), or that any analysis capable of capturing them must contain too many *ad hoc* epicycles to express our concept or to capture the real relation we conceive as knowledge. Neither claim precedes Nozick's willingness to dismiss such convictions as that I know that I am not a brain in a vat, that my son is not an assassin, that he is not a cleverly constructed robot, and so on. Such intuitions (especially the more local ones) do not seem particularly vulnerable. Indeed they seem as certain as any convictions regarding particular claims to empirical knowledge from which I might begin to reconstruct my concept. Thus they reflect badly on an account that must dismiss them.

Perhaps more serious epistemologically than particular counterintuitive implications of the analysis regarding more local skeptical alternatives to our everyday beliefs is the incompleteness of the more general debate with the skeptic. Any analysis that avoids Gettier examples by specifying a nonaccidental connection between the knower's belief and the fact to which it refers can provide a partial answer to the skeptic. Its advocate need only stipulate that the connection in ques-

tion can hold in some circumstances but not in others that are indistinguishable from the former "from the inside," from the perspective of the knower's experiential states.<sup>8</sup> If the connection alone is what counts, and not how things appear to the knowing subject, then knowledge is not threatened by skeptical alternatives that appear indistinguishable from the internal perspective. Nozick's claim that knowledge obtains under tracking—since the skeptic's worlds (of brains in vats and so forth) are not the closest in which our beliefs are counterfactually false—is but a special case of this more general play open to an advocate of this kind of analysis.

But if the possibility of knowledge is saved in this way, there is nevertheless no *demonstration* of knowledge in the face of skeptical challenge here. The epistemologist's response to the skeptic is seriously incomplete. We are left pretty much with a draw, in which, *for all we know*, the skeptic's worlds are as possible as the real world in which we believe. Nozick simply states or assumes that the skeptical alternatives are not the closest worlds in which our everyday beliefs are false. If they are indeed very distant from the actual world, then of course they are not as possible as the real world in which we believe. But his assumption is entirely unsupported and indeed must remain so according to his analysis. If we cannot know that the skeptic's worlds are not actual (as on his account we cannot), then we surely cannot know that they are distant from the actual world.

It might be thought that the draw can be avoided by moving from claims about knowledge to considerations of rational belief and evidence.<sup>9</sup> Nozick in fact points out that there are asymmetries between the evidence for the existence of the real world and the evidence for the existence of the skeptic's worlds (264). There is evidence for the real world, since, if there were no such world, our experience would not be as it is; while there is no evidence for the skeptic's worlds, since, if they did not exist, our experience would be exactly as it is. Our experience as evidence tracks the hypothesis that there is a material world, but it does not track the skeptical alternatives. Thus, if rational belief is to be based on the evidence, it might appear more rational to believe in the real world than in the skeptical alternatives. This outcome might suffice to grant the victory to the nonskeptical epistemologist in his debate with the skeptic.

Winning this particular round (if indeed we have) is at the same time losing the fight, however. The problem is that the evidential relation is defined by Nozick in counterfactual terms similar to those in the analysis of knowledge. *e* is (strong) evidence for *h* if it tracks *h*. According to Nozick, our experience tracks the real world but not the skeptic's worlds. But since the evidential relation is an external to our point of view as the knowledge relation itself, we again have no way of knowing

this. We have access to the supposed evidence, but not to the evidential relation. In the end the latter must simply be stipulated or assumed to obtain. Thus we are left again with only the possibility of knowledge, not with anything approximating to its demonstration.

Nozick may respond that refutation of the skeptic is not among his purposes, or that it is indeed impossible. For him the lure of the skeptic's arguments seems to lie in this very impossibility of refutation. But there is an alternative explanation, and one that will satisfy the true believer more. This explanation points to the link between skepticism and the tendency to make the conditions for knowledge too strong, a link with a long tradition in epistemology. The stronger the analysis of knowledge (and of evidence), the easier the skeptic's task. We saw above that Nozick's analysis appears too strong to capture our concept of knowledge as reflected in convictions regarding particular cases. It therefore may well fit the tradition that renders skepticism more plausible and immune to refutation than it need be.

### III

The problems here specified for the counterfactual analysis should motivate us to seek an alternative, despite the extent of the literature on this subject in recent years. (Measured against other areas of philosophical inquiry, progress in meta-epistemology has been actually rapid.) In so doing we may either retain Nozick's approach or opt for one of the other strategies for accommodating the now extensive file of examples and variations. In my view Nozick's approach, seeking the right connection between beliefs and facts, is entirely correct. Defeasibility and reliability accounts (that are truly distinct) face more insuperable difficulties.

The difficulties for accounts that require nondefeasible justification—justification that cannot be overturned by evidence not in the subject's possession—are of two kinds. First, nondefeasible justification implies justification, and in my view this is already too strong, at least for ordinary (internal) notions of justification. Examples that show this share a common form. They involve believers whose beliefs connect to facts in required ways but who do not know that their beliefs are so connected and who are not justified (from their points of view) in thinking so. A medium who suddenly acquired the power to truly predict the future and whose predictions begin to turn out always right would have to be said to know the future. She would know it from the time that her predictions began to be right, even though at that time she was not justified in her beliefs. Space prevents a complete defense here of the controversial import of such examples.

Less controversial is the demand that such accounts be able to sepa-

rate misleading evidence that defeats justification from that which may be safely ignored. Not all misleading evidence defeats justification. If I know that Haveit owns a Ford by having the usual kinds of evidence, then the fact that Haveit has an enemy in some other city who might convince me in that distant possible world in which I meet him that Haveit stole the car he always drives is irrelevant to my present knowledge. The usual strategy for defeasibility analyses is to consider whether justification survives the correction of a subject's false beliefs. Omniscience cannot be granted in the counterfactual situation, however, since that would entail that my knowledge can never be defeated by evidence I do not possess. But suppose in our example that Haveit's enemy is in my own office, and I just happen not to have heard the false rumors that have convinced everyone else in the office.<sup>10</sup> The usual ploy for the defeasibility advocate is to correct only my actual false beliefs.<sup>11</sup> But this may not work here either, since my knowledge can then be saved simply because I have not formed certain beliefs that I should have.

Consider finally the following example here. Suppose that I have (misleading) evidence that Nogot owns the Ford he drove to work this morning, but I am very cautious. I therefore believe the disjunction: either Nogot owns a Ford or he does not own the blue Ford he drove to work this morning. The real situation is that Nogot rented the Ford he was driving but that, unknown to him (or me), he has won that very car in the rental company's annual lucky giveaway lottery.<sup>12</sup> Here I do not know the disjunction to be true, since I know neither disjunct (and there is a close world in which he has won a different Ford). Yet it seems that any available defeater for my belief in the first disjunct would justify me in believing the second, and conversely. Thus I would continue to believe truly and be justified in believing when seeming defeaters are added. The true propositions of which I am unaware would not prevent me from being justified in believing (truly) the disjunction, and yet I do not know it. I do not know how to handle such examples within the defeasibility framework, and I suspect that any condition that could handle them would be so complex as to render the resulting analysis unusable as a guide for normative epistemology.

As for reliability accounts, these differ according to how the concept of reliability is unpacked. If sources for beliefs are required to be only generally reliable, then the account is too weak. This becomes clear in lottery examples. We can make the reliability of the inductive source of my belief that my ticket will not win a lottery as close to 1 as we like by increasing the number of tickets. Yet no matter how high the number, I still do not know in advance that my ticket will not win. (If I knew this in advance, then it would be irrational for me to buy a ticket. But it may not be irrational to enter the lottery—it depends on the odds.) Nor

can we demand perfect reliability from a source in its general operation, for that demand would disallow all knowledge gained from perception or inductive inference. We must therefore think of the sources for belief only as they operate in particular contexts. A source of belief is reliable in context if and only if it produces true beliefs all or most of the time in such contexts. Given the source of a belief in a particular situation, facts of the same type as that to which the belief refers must tend to obtain.

Properly elaborated, this requirement singles out a particular kind of (indirect) connection between fact and belief that turns out to be an instance of a more general relation that I shall later endorse as a genuine constraint on knowledge. It should be clear at this point, however, that we have returned to the Nozickian approach of seeking the proper connection of belief to fact. Reliability in this sense has little to do with justification traditionally construed (as a relation between a subject or his belief and his accessible evidence). We may then turn again directly to the quest for the right connection. Recall that we require a relation weaker in itself than Nozick's, but with an additional constraint on connecting chains that renders it also stronger than his in some contexts.

Nozick actually guides us further toward our goal in his discussion of evidential support for a belief or hypothesis. The support lent by a body of evidence to a hypothesis is measured by the probability of the evidence given the hypothesis less the probability of the evidence given the falsity of the hypothesis:  $\text{pr}(e/h) - \text{pr}(e/\text{not-}h)$ . If we seek a connection of belief to fact weaker than the counterfactual tracking relation, we might try requiring a relation similar to this notion of support. We could say that the probability of a belief given the fact to which it refers, given its truth, must be greater than the probability of the belief given its falsity:  $\text{pr}(\text{belief}/p) - \text{pr}(\text{belief}/\text{not-}p) > 0$ . We can be more stringent by raising the degree to which these probabilities must differ.

Nozick calls the relation just described (when it involves inference from evidence) "almost-knowledge." Can we simply accept it as the knowledge relation? Not quite. We saw above that when a proposition believed is false only in a very distant possible world, it is irrelevant to a knowledge claim whether it would be believed in that world. Likewise, the probability of belief in that world is irrelevant. I know that my son is neither an assassin nor a cleverly constructed robot, however probable or improbable are my beliefs to the contrary in those distant (but closest) worlds in which my present beliefs are false. Thus we do not need to subtract the second probability from the first in the formula proposed above as expressing the knowledge relation.

In my earlier paper on the analysis of knowledge I spoke rather of the relevant fact's significantly raising the antecedent probability of



the belief's being held:  $\text{pr}(\text{belief}/p) > \text{pr}(\text{belief})$ . When this relation holds across close possible worlds (when the probability relations are lawlike), then the truth of the belief, or the fact to which it refers, prominently helps to explain why the belief is held. This explanatory relation is the connection necessary for knowledge.<sup>13</sup> While it does not require us to consider distant possible worlds, and is therefore weaker than tracking, it must be combined with an additional constraint on explanatory chains. Each (probabilistic) link in such a chain must itself make the next link more probable. This idea involves second order probabilities of probability relations (and therefore galaxies of possible worlds under that interpretation). It is therefore somewhat complex to conceive and apply, yet intuitive as a constraint on explanatory chains. It determines when prior explanations in such chains are relevant to later ones. These conditions together handle the counterexamples proposed above both to Nozick's analysis and to the other major types of account.

#### IV

We may briefly review those examples to see how this alternative connection holds up in them. In the assassination example, there exists an explanatory chain that accounts for my knowledge of my son's innocence. Facts that explain his innocence — that someone else commits the murder, that my son is playing tennis with me at the time, and so on — also explain my belief in his innocence. Someone else's committing the murder explains the radio broadcast, which explains my having any beliefs about the matter at all. Thus a fact that explains my son's innocence can help to explain my belief about it, even though I would not believe him guilty if he were. A similar chain exists in the thermometer example. The fact that it is 70 degrees explains the thermometer's reading, which explains my belief that it is not — 50 degrees. That fact that it is 70 degrees, or facts that explain why it is 70 degrees, also explain why it is not — 50 degrees.

Consider next the intentionally set up Gettier case in which Haveit deceives me into truly believing that someone in my office owns a Ford. Here there is an explanatory chain in which the fact that someone does own a Ford (namely Haveit), together with Haveit's intention to deceive me, explains his presenting me with the deceptive evidence, which in turn explains my true belief. But the constraint on explanatory chains is violated. My believing on the basis of the evidence presented is not made more probable (explained) by the evidence's being connected with the fact that makes my belief true. A similar violation occurs in the case of the brain in the vat programmed to believe that it is a brain in a vat. The truth of the belief helps to explain why the pro-

grammer programs it, but that connection does not help to explain why the brain believes on the basis of being programmed. Contrast these cases with ordinary perceptual beliefs, where our believing on the basis of sensory evidence is itself made more probable (given natural selection of this type of belief) by the usual connections between perceived facts and sensory evidence.

Next came the example of the logician with failing capacities, who still knew a particular conclusion by having inferred it from true premises. Here the explanatory analysis differs from Nozick's in being again weaker. The connection between fact and belief need not hold in all close possible worlds, as in the counterfactual account. It suffices for the truth of the belief to raise the probability of its being held that there are more close worlds in which the two are connected than worlds in which they are not. In this example, the fact that there is one close world in which the logician's inferential capacities have failed and he no longer believes the conclusion does not defeat his claim to knowledge in the actual world, so long as there are other close worlds in which other factors change yet he continues to infer correctly.

Epistemically more important examples pertain to both local and general skeptical alternatives to our everyday beliefs. We saw that Nozick's analysis precludes knowledge that such alternatives fail to obtain, and that this in turn precludes any satisfactory answer to the skeptic's challenge, any way to show that we have everyday knowledge in the face of that challenge. Knowledge remains possible but non-demonstrable according to Nozick's analysis. Using the explanatory account, the epistemologist can do better. He can show that we have knowledge by showing that explanations for our beliefs in common sense or scientific terms are better than the skeptic's alternative explanations. According to the ordinary inductive criteria for good explanations, the former certainly are better (simpler, more plausible, deeper, less ad hoc, generating fewer unanswerable additional questions, and so on). The difficult part of this task is showing that our ordinary and scientific inductive criteria tend to be truth preserving, that they lead us or connect us to real explanatory connections among phenomena or objects in the world. I have attempted to provide this difficult argument elsewhere and cannot begin to repeat it here. However plausible my version of it is, it is clear that our prospect for answering the skeptic is better given the explanatory analysis than under Nozick's, which in itself blocks any satisfactory answer.

The counterexamples to defeasibility and certain kinds of reliability analyses can be disposed of briefly. Consider again the two variants of the case in which I claim to know that Haveit owns a Ford. In the first Haveit has an enemy in some distant city who convinces people to believe nothing that Haveit says, while in the second that enemy is in my

office and has (falsely) convinced everyone else in the office that Haveit stole the car he is driving. In the first case what explains my belief is the evidence I have received, which in turn is linked to the fact that Haveit does own the Ford. The distant false rumors are irrelevant to my knowledge claim, since the possible world in which I learn of them is also distant, and we are interested in probability relations only across close worlds. But in the second case, if there are many close worlds in which I hear the false rumors, then my belief is explained not simply by the evidence I possess, but by my happening not to have heard the rumors, since the evidence raises the probability of my belief only in the context of my not having heard them. Thus our intuitions are supported by the analysis in both cases.

The disjunction case that is so problematic for defeasibility accounts presents no problem here either. My lack of knowledge (despite remaining justified in my true belief when supposed defeaters are added) is explained by the lack of any explanatory or probabilistic connection between the fact that Nogot has just won the car and my belief in the disjunction.

Finally, the explanatory analysis correctly evaluates the lottery example that defeats straightforward fallibilistic reliability accounts. The best explanation for my not winning involves the drawing of another ticket (which in turn is explained by the configuration of tickets in the barrel, the position of the person drawing the winner from the barrel, and so on — not simply by the number of tickets). But these facts are unconnected with my belief prior to the drawing that I will not win. Nor are they (sufficiently) explained by what explains my belief — that is, the large number of tickets. (Remember that in most contexts we require that the fact significantly or prominently raise the probability of the belief's being held. Contexts vary in this respect, however, and perhaps if we wanted to be extremely lax here, we might grant knowledge prior to the drawing. The point is that the reliability account must grant it, whereas ordinarily we would not.) The more promising way of spelling out the notion of reliability turns out to instantiate one type of permissible explanatory chain, in which the source of belief explains both the belief and the fact to which it refers (or it is explained by that fact).

## V

In explaining the significance of knowledge for knowing subjects, Nozick argues that evolution or natural selection could be expected to result in subjects not simply with (many) true beliefs, but with sensitivity of beliefs to (changing) facts (284). He takes this sensitivity to consist in the ability to track facts. Thus he takes evolutionary theory

to support his analysis, which helps to show how the ability to know is an important achievement for the species as well as for the individual. Knowledge as tracking is the type of cognitive state that might be selected (or rather the ability to achieve such states might be selected). An analysis of knowledge should suggest why individuals and species benefit from its acquisition, and Nozick claims that his account satisfies this desideratum.

But in fact tracking is not quite the relation we would expect between believer and fact as the result of natural selection. We would not expect or require this relation to propositions, if many propositions are false only in distant possible worlds. Survival is not enhanced by sensitivity to counterfactual occurrences in those worlds. In fact, as we saw from several earlier examples, sensitivity to such distant counterfactual situations might distract us from knowledge in the real world. On the other hand, facts must be (probabilistically) connected with beliefs across close possible worlds and not simply in the actual world. We must be sensitive to likely changes in empirical conditions that might falsify our beliefs in close possible worlds — that is, those that might well become actual. From this theoretical viewpoint too we require a connection weaker than Nozick's tracking, which takes us to distant worlds in many cases.

Nozick notes that knowledge as a significant cognitive state for subjects should fall somewhere between having beliefs that just happen to be true and having beliefs that vary with the facts to which they refer in all possible worlds (285–86). The reason why the latter is not required, however, is precisely that we need not connect to variations in facts in distant possible worlds, in circumstances very unlike our actual environments that are also very unlikely to occur. But his tracking requirement may demand just such connections when considering the falsity of our present beliefs. Nozick here suggests a fruitful criterion for a successful account of the connection between beliefs and facts that constitute knowledge. He also refers us to a spectrum of accounts differing in the strength of the modal relation they require. In my view he missed finding the right connection by just a short distance within that spectrum.

## NOTES

1. E. Gettier, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121–23.
2. The classic attempt was Alvin Goldman, "A Causal Theory of Knowing," *Journal of Philosophy* 64 (1967): 357–72.
3. See Keith Lehrer, *Knowledge* (Oxford: Oxford University Press, 1974); Marshall Swain, *Reasons and Knowledge* (Ithaca: Cornell University Press, 1981).
4. See, for example, Alvin Goldman, "What Is Justified Belief?" in *Justification and Knowledge*, ed. G. Pappas (Dordrecht: Reidel, 1970): 1–25.

5. Robert Nozick, *Philosophical Explanations* (Cambridge: Harvard University Press, 1981), ch. 3; see also Fred Dretske, "Conclusive Reasons," *Australasian Journal of Philosophy* 49 (1971): 1-22.
6. Parenthetical references are to Nozick's book *Philosophical Explanations*, op. cit.
7. Similar problems for Nozick's account are raised by Robert Shope, "Cognitive Abilities, Conditionals, and Knowledge: A Response to Nozick," *Journal of Philosophy* 81 (1984): 29-48.
8. Compare Steven Luper-Foy, "What Skeptics Don't Know Refutes Them," *Pacific Philosophical Quarterly* 65 (1984): 86-96.
9. This move was suggested to me by Frederick Schmitt.
10. The example is now like one from Gilbert Harman, *Thought* (Princeton: Princeton University Press, 1974), 143-44.
11. Lehrer, op. cit., chap. 9.
12. This is a variant of an example in Robert Shope, *The Analysis of Knowing* (Princeton: Princeton University Press, 1983), 216.
13. Alan Goldman, "An Explanatory Analysis of Knowledge," *American Philosophical Quarterly* 21 (1984): 101-8.

## SIX

# TRACKING, CLOSURE, AND INDUCTIVE KNOWLEDGE

Jonathan Vogel

## I. INTRODUCTION

One of Robert Nozick's chief concerns in *Philosophical Explanations* is to refute a classical argument for Cartesian skepticism. According to the skeptic, we don't know that we aren't victims of a massive illusion or deception. For that reason, we have little or no knowledge about the way the external world really is. Nozick concedes the first point, that we don't know that we aren't systematically deluded, but he rejects the legitimacy of the inference which takes the skeptic to his conclusion.

Nozick believes that, in making this inference, the skeptic at least tacitly assumes that knowledge is closed under known logical implication. In other words, the skeptic assumes that the following epistemic principle is valid:

Closure Principle: Where *S* is a subject, and *p* and *q* are propositions, if *S* knows that *p* and knows that *p* entails *q*, then *S* knows that *q*.<sup>1</sup>

Nozick argues that this principle doesn't hold, and that the skeptical argument which depends upon it doesn't go through. His procedure is to present and defend an analysis of knowledge that explicitly provides for closure failures. If that analysis is really correct, then the Closure Principle doesn't hold, at least in its full generality.

A key feature of Nozick's analysis is a condition that he calls the tracking requirement. In this essay, I will examine the role of the tracking requirement in Nozick's analysis and attempt to clarify the relation between tracking and the Closure Principle. I will consider various problems with the notion of tracking, and argue that the tracking requirement is incompatible with a satisfactory account of in-