

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 .¹

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-

ductive knowledge. My conclusion is that Nozick doesn't succeed in refuting the skeptic after all.

II. NOZICK'S TRACKING CONDITION

According to Nozick's distinctive account of knowledge, a knowing subject must satisfy an adherence condition (A) and a tracking condition (T). These requirements may be spelled out in a preliminary way as follows:

(A) If p were true, S would believe that p

(T) If p were false, S would not believe that p .

In this essay, I will be concerned with the tracking condition (T), and I will have little to say about (A) as such.

Nozick intends (A) and (T) to be read as standard English subjunctive conditionals. The most widely accepted theory of such conditionals analyzes their semantic properties in terms of a space of possible worlds, in which worlds are closer or farther from the actual world to the extent that they are more or less similar to that world. According to this analysis, due principally to David Lewis and Robert Stalnaker, a subjunctive conditional of the form $(t \rightarrow v)$ is true just in case v holds in the closest world (or worlds) to the actual world in which t is true.² Thus, (T) in particular says that S knows that p only if S doesn't believe that p in the closest possible world(s) in which not- p holds — that is, the closest possible world(s) in which p is false.³

However, (T) as stated cannot stand; there are clear cases of knowledge in which the appropriate instance of (T) is not satisfied. Alvin Goldman has illustrated this point with a nice example:

Condition (3) also doesn't seem necessary for knowing. Consider this example. Oscar sees Dack the dachshund and believes there is a dog before him. If there weren't a dog before him, a hyena would be there instead, which Oscar would misclassify as a dog. Letting p = 'There is a dog before me,' condition (3) is violated; but Oscar still knows.⁴

To meet this difficulty and others like it, Nozick revises his original formulations to take into account the method by which a subject arrives at his belief. (T) is replaced by

(T') If p were false, it would not be the case that S would believe that p via method M (where M is the method S actually uses)⁵

This revised analysis escapes Goldman's "Dack the Dachshund" example. The problem for the original version was that, although Oscar knew there was a dog before him, he did not satisfy (T). But Oscar does satisfy (T'), on a natural reading of it. (T') requires, in this case, that if

does not!

there
metho
car's a
squat
etc., c
dog b
ever,
mal, v
clude
metho
he wc
(T') a
Th
ject's
meth
diate
well l
the o
any u
for N
issue
and t
ment

No

tre
th
m
pl
fe
in

Let u
that
base
actu
factu
the l
the s
Th
the s
exac
that
were

there weren't a dog before Oscar, he wouldn't believe by *his actual method* that there was a dog before him. One way of characterizing Oscar's actual method would be something like 'If you seem to see a small, squat, elongated animal with short legs and floppy ears, that barks, etc., conclude that there is a dog before you.' Now, if there weren't a dog before Oscar, he would as it happens, believe that there was. However, his method would be roughly, 'If you seem to see a spotted animal, with short back legs and a bristly mane, that has a shrill cry, conclude that there is a dog in front of you.' This is not Oscar's actual method. So, it's the case that: if there weren't a dog in front of Oscar, he wouldn't believe by his *actual method* that there was. Oscar satisfied (T') after all.

This example points up how important the specification of a subject's method is to the success of Nozick's analysis. The concept of the method by which a person comes to believe something is not an immediately clear one; in the "Dack the Dachshund" example, there may well be better or more natural ways to describe Oscar's methods than the ones I gave. In fact, serious doubts have been raised as to whether any uniform construal of the notion of a method of knowing is suitable for Nozick's general purposes.⁶ However, no resolution of this larger issue is needed in order to evaluate Nozick's views about skepticism and the Closure Principle. Rather, the following, more limited treatment is sufficient.

Nozick observes at one point that his view

treat[s] the method he [the subject] uses as identified from the inside, to the extent that it is guided by internal cues and appearances. . . . The method used must be specified as having a certain generality if it is to play the appropriate role in subjunctives. This generality is set by the differences the person would notice; the methods are individuated from the inside. (PE, 232-33)

Let us say, then, that for (T') it is a *sufficient* condition for S's believing that *p* by the same method as S's actual one that: S's belief that *p* is based on exactly the same evidence and inference patterns (if any) as S actually utilized.⁷ In other words, if it were the case, in some counterfactual situation, that S had the same evidence *e*, and believed that *p* on the basis of *e*, in that counterfactual situation S would believe that *p* by the same method S actually used.

The idea of (T') is that if *p* were false, S would not believe that *p* by the actual method S used to arrive at *p*. We just observed that if S had exactly the same evidence and believed *p* on that basis, S would believe that *p* by his actual method. (T'), then, requires as a *minimum*, that if *p* were false, S would not believe that *p* on the basis of his actual evi-

PART
t No-
OWL-
dence
dence
as D.
ge S.
Fu-
the
king,
athan
SKEP-
even
cism,
Peter
on F.
ism,
the
es
stant
sity,
rna-
om-

dence. This requirement is quite weak, since it is satisfied so long as *S* would have had different evidence, had *p* been false—regardless of what *S* believes in that situation. From now on, let (T)' mean this minimal condition.

III. NOZICK ON SKEPTICISM AND CLOSURE

It follows directly from Nozick's account of knowledge that none of us knows the various skeptical hypotheses to be false. That is, no one knows not-SK, where not-SK says of that person that he is not, for example, a systematically deceived brain in a vat. The reason why we do not know not-SK, on Nozick's account, is that each of us fails to satisfy the appropriate instance of (T)'. For Nozick himself, the instance of (T)' would run:

- (1) Not-(not-SK) → not-(Nozick believes, on the basis of his actual evidence, that not-SK)

Less schematically:

- (2) If Nozick were a brain in a vat, he wouldn't believe, on the basis of his actual evidence, that he wasn't a brain in a vat.

Surely (2) is false. If Nozick were a brain in a vat, the course of his experience would have been, by hypothesis, exactly the same as it actually has been. So, Nozick would have had the same beliefs and evidence he actually has, including the belief that he is not a brain in a vat. But (T)' requires that he not have that belief under the circumstances where he is a brain in a vat. So, (T)' isn't satisfied, and Nozick does not, by his theory, know that not-SK. He does not know that he isn't a brain in a vat.

Although the tracking requirement excludes knowledge that the skeptical hypothesis is false, it doesn't impair knowledge of various particular facts about the world. The example Nozick himself uses is that he still knows he is sitting in a chair in Jerusalem as he writes his book (206–7). The relevant instance of (T)' is:

- (3) If Nozick weren't sitting in a chair in Jerusalem writing his book, he wouldn't believe, by his actual evidence, that he was.

(3) is true. If Nozick weren't sitting in a chair in Jerusalem, he would have been doing something else in Jerusalem, or perhaps gone somewhere other than Jerusalem. But in these counterfactual situations, he would have remained aware of where he was and what he was doing. So, he would not have held the erroneous belief that he was sitting in Jerusalem when he wasn't. A fortiori, he wouldn't have believed by his actual evidence that he was sitting in Jerusalem when he wasn't. In this

Tracking

way, N
ment,

The:
tions
scribed
Jerusa
temati
Nozick
saw, o
not kn
ure is
under
(J ⇒

Now
known
are the
edge is
satisfy
other f
sure of

How
problem
conditi
that in
p—tha
K(2) S
K(1) ar
identical
cal imp
'contra
track' t
'Contra
though

It is e
sition.
ing its
edge ca
p, S m
ones. S
dition f
quence
that he
erally n
for the

way, Nozick's belief about where he is does satisfy the tracking requirement, and is something he may know.

These examples show how the tracking requirement leads to violations of the Closure Principle for knowledge. In the situation described, Nozick knows that (J = 'Nozick is sitting in a chair in Jerusalem writing his book') entails (not-SK = 'Nozick isn't a systematically deceived brain in a vat near Alpha Centauri'). That is, for Nozick, J has not-SK as its clear logical consequence. But, as we just saw, on Nozick's analysis of knowledge, he may know that J while not knowing that not-SK.⁸ Thus, the Closure Principle fails. This failure is due to the fact that the tracking condition itself isn't closed under known logical implication. For, despite Nozick's knowing that ($J \Rightarrow$ not-SK), he satisfies (T) for J but not for not-SK.

Now it seems reasonable to think that knowledge is closed under known logical implication only if its component necessary conditions are themselves so closed. If that is so, Nozick can establish that knowledge isn't closed just by showing that, in order to know, one must satisfy the tracking requirement. Questions arising in connection with other features of Nozick's analysis become irrelevant, so far as the closure of knowledge and skepticism are concerned.

However, the assumption on which this simplification depends is problematic. The view that knowledge is closed only if its component conditions are faces the following sort of counterexample.⁹ Suppose that in order for S to know that p it is required that: K(1) S tracks p —that is, S and p satisfy the appropriate instance of (T) or (T)'; and K(2) S tracks all of p 's logical consequences. Call the satisfaction of K(1) and K(2) 'contracting.' For the reasons just considered K(1) as identical to either (T) or (T)' is not closed, by itself, under known logical implication. But the point is well taken. If S 'contracts' p , S will 'contract' all of p 's logical consequences, and therefore S will 'contract' the subset of those logical consequences that are clear to S . 'Contracting,' then, is closed under known logical implication, even though one of its component conditions is not.

It is easy to imagine the following reply from friends of Nozick's position. If knowledge were 'contracting,' it could be closed without having its component conditions severally subject to closure. But knowledge can't be 'contracting' or anything entailing it. For if S 'contracts' p , S must 'contract' *all* of p 's logical consequences, not just p 's clear ones. So, if knowledge were 'contracting,' it would be a necessary condition for knowledge of p that a person know all of p 's logical consequences. But no one can know all the logical consequences of anything that he knows; so, if knowledge were 'contracting' one could know literally nothing. And that sets the standard of knowledge too high even for the Cartesian skeptic.¹⁰

PART
t No-
OWL-
dence
dence
as D.
ge S.
A. Fu-
g the
cking,
athan
SKEP-
teven
cism,
Peter
en F.
cism,
t the
es
istant
rsity,
erna-
com-

Another point worth noting is that, if closure holds despite the tracking requirement, Cartesian skepticism is inescapable. For, as the previous discussion indicates, belief in the falsity of the skeptical hypothesis fails to satisfy Nozick's tracking condition; thus, no one knows that the skeptical hypothesis is false. However, by the Closure Principle, we have to know that the skeptical hypothesis is false in order to know virtually anything about the external world. It follows, then, that we will have very little knowledge of the world if both the tracking condition and Closure Principle are valid. This outcome would be unwelcome to anyone who finds the Closure Principle intuitive and skepticism incredible.

It is clear, then, that the attempt to defend closure while granting that (TY) may be necessary for knowledge is an approach which faces serious liabilities. An alternative line of thought, which I shall pursue, is to try to show that (TY) is in fact *not* necessary to knowledge. If that is correct, then it is irrelevant, for my purposes, whether knowledge can be closed even if its component conditions are not.

IV. SOME PROPOSED COUNTEREXAMPLES

For the reasons just given we may grant the assumption that, if (TY) is a necessary condition for knowledge, knowledge itself doesn't obey the Closure Principle. Such a result provides an important perspective on the large literature which has appeared in response to Nozick's book. Any number of counterexamples have been offered to refute Nozick's analysis of knowledge. Although many of them are controversial, some, I think, succeed. However, almost every counterexample offered aims to show that Nozick's analysis fails by being too *weak*; that is, the analysis is supposed to be defective because it fails to exclude cases which it ought not to count as knowledge. It is important to note, though, that counterexamples of that type do not undercut Nozick's case against closure, since they do not show that (TY) is not necessary for knowledge.

Robert Shope, however, has argued against the validity of the tracking requirement. Shope maintains that some subject S can know that (C): 'It is true of some of those beliefs that S does have about beliefs that they might be otherwise; that is, that S might not have them.'¹¹ The tracking requirement would have us consider counterfactual situations in which (C) is false. Here, though, the limitation to beliefs about beliefs is cumbersome and implausible. If first-order beliefs could vary while second-order beliefs were fixed, in many cases S could believe that he believed that *d*, but not believe that *d*; also, S could believe that *f*, without believing that he believed it. It is problematic to characterize anyone subject to widespread incongruities like

these a
grasp.
true of
that S

Shope
cannot
(C) we
dence.
S's beli
pends.
dence,
we can
knowle

Ther
beliefs
be an e
normal
be. Sho
ternal p
For (C)
beliefs
false.¹³
decepti
neither

Actu
membe
Nozick
in some
beliefs.
to know
that I a
believe
that the
count.
none th
satisfy
I would
my act
claim b

Pres
oursel
regard
tiveness
tions is

these as having beliefs at all, so Shope's example becomes difficult to grasp. For Shope's purposes, (C)' would be better, where (C)' is: 'It is true of some of the beliefs S has that they might be otherwise; that is, that S might not have them.

Shope's point can now be put this way: S may know that (C)' yet S cannot satisfy (T)' for (C)'. The appropriate instance of (T)' says that if (C)' were false, S would not believe (C)' on the basis of S's actual evidence. But if (C)' were false, all of S's beliefs would be fixed, including S's belief that (C)' and all the beliefs, if any, upon which that belief depends. So, if (C)' were false, S *would* believe that (C)' on his actual evidence, making (T)' *false*. Yet, it seems that (C)' holds is the sort of thing we can and do know. So, (T)' cannot be a necessary condition for knowledge.

There are still some mysteries about this example. It is clear that if S's beliefs were such that they could not have been otherwise, S wouldn't be an embodied, mobile person, with functioning senses, negotiating a normal environment—although that is what S would take himself to be. Shope suggests that "one can at least imagine a time at which an external power has forced on S those beliefs which S has about beliefs."¹² For (C)' to be false, the "external power" would have to force all of S's beliefs upon him, and a great many of those beliefs would have to be false.¹³ In short, (C)' is the denial of a skeptical hypothesis of massive deception, and it would be perfectly consistent for Nozick to say that neither S nor anyone else knows (C)'.¹⁴

Actually, Shope's example can be seen as an especially baroque member of a family of examples which might be deployed against Nozick's analysis. If S knows that (C)', S knows that he is not deceived in some crucial respects about the processes by which he arrives at his beliefs. On Nozick's analysis, it appears that it is impossible for anyone to know that he is not deceived about anything in particular. Imagine that I am looking at the statement I just received from the bank, and I believe my balance to be as it appears on the statement. I also believe that the bank is not deceiving me about how much money is in my account. Do I know that the bank isn't tricking me? If it were, I would be none the wiser, since by assumption I would be taken in. So, I do not satisfy (T)' in this case; it is not true that if the bank were deceiving me, I would not believe that the bank was not deceiving me (on the basis of my actual evidence). The same line of thought can be applied to any claim by someone to know that he is not deceived in some way or other.

Presumably, Nozick would accept the result that we do not know ourselves to be undecieved by banks or anyone else. In effect, he would regard these as small-scale skeptical hypotheses. The counterintuitiveness of our not knowing ourselves not to be victims of small deceptions is mitigated for Nozick by the failure of the Closure Principle in

these cases. For instance, suppose that, in the example just given, the bank is really honest. Even if I do not know that the bank is not tricking me, I do know what my balance is. If my balance had not been what it actually is, say \$100, the bank, being honest, would have reported the different figure and I would not have believed that my balance was \$100. In other words, I can satisfy (T)' for my belief that my balance is \$100, without satisfying it for the belief that the bank is not deceiving me. Consequently, I can know what my balance is, without knowing that the bank isn't deceiving me about my balance. This situation is analogous to the circumstances created by grand skeptical hypotheses postulating near-total deception.¹⁵

The possibility of defending (T)' in this way raises a general methodological problem. Suppose an apparent counterexample to the tracking requirement is found, in which someone seems to know that *p* even though *S* doesn't satisfy (T)' with respect to *p*. To make the counterexample legitimate, must it also be shown that *p* isn't the denial of what of what is, in some *extended* sense, a skeptical hypothesis? It is difficult to see what such a demand would really come to, or how it could be met. Still, as I will suggest, there are counterexamples to the tracking condition which seem to obey any reasonable version of this constraint.¹⁶

V. NON-CLOSURE AND THE PROBLEM OF INDUCTIVE KNOWLEDGE

I would like at this point to turn to Nozick's treatment of inductive knowledge. By inductive knowledge, I mean knowledge of some proposition on the basis of a limited sample, such that what one knows about the sample does not entail the truth of the proposition known. Inductive knowledge may be of general propositions (for example, 'All pure samples of silver are good conductors'). There is also inductive knowledge of singular propositions (such as, 'The next emerald you see will be green'). Nozick believes that inductive knowledge can be handled by his account to the extent that all genuine cases of inductive knowledge satisfy (T)':

Do we know that the sun will rise tomorrow, that the earth will continue to rotate its axis during the next 24-hour period? If the sun were not going to rise tomorrow, would we have seen that coming, would that alteration in the earth's rotation have been presaged in the facts available to us today and before? If so, then we do know the sun will rise tomorrow; our belief that it will tracks the fact that it will, by being based on facts that would have been different otherwise. But isn't it logically possible that everything was as it was until now, yet the earth will not continue to rotate tomorrow? Yes, there are such skeptical logical possibilities SK: the bread no longer nourishes us, the sun stops in the sky, an event of a

cert
elal
hav
ske
do
ind
pre

This :
Nozic
will co
spin,
cause
in this
would
earth
earth'
it; we
will k
lieve
false,

As
know
at all,
ning :
anom
Closu
spinn
won't
know
Even
won't
then,
(T)' d

Mc
has tl
ositic
wise,
hence
of sk
be, tl
track
may
prop
spinr

certain sort no longer continues to produce its usual effects. If they are elaborated suitably, so that everything we can detect up until now would have remained the same, then we don't know they do not hold. The skeptic about induction is right to say we don't know these possibilities do not hold, but he is wrong to deny we know those particular results of inductive inference whose falsity would have been reflected back and presaged in the facts upon which we base the inference. (PE, 222-223).

This approach leads directly to an odd consequence. Suppose that Nozick is right, and we can know in the way he describes that the earth will continue to spin. While we do know that the earth will continue to spin, we do *not* know that it won't *suddenly* stop spinning. That's because the relevant instance of (T)' is false for the latter. That condition in this case reads 'If the earth were suddenly to stop spinning, you (or I) would not have believed on the basis of our actual evidence that the earth would not suddenly stop spinning.' But, if the change in the earth's rotation were to be sudden, we would have had no warning of it; we would have the same expectation we always have that the earth will keep going without stop, sudden or otherwise. Hence we would believe that 'The earth will not stop spinning suddenly' even if it were false, and thereby fail to meet the condition (T)'.

As the above quotation indicates, Nozick seems prepared to acknowledge that someone can know that 'The earth won't stop spinning at all,' yet, at the same time, not know that 'The earth won't stop spinning suddenly.' We could expect Nozick to explain the impression of anomaly here as due, once again, to an unreasoned commitment to the Closure Principle. For, it may be that we believe the earth won't stop spinning suddenly because we infer that from the wider belief that it won't stop spinning at all. In that case, the reason for thinking that we know the former would have to be some form of the Closure Principle. Even so, Nozick's treatment of this example is inadequate, if 'The earth won't stop spinning suddenly' is the sort of thing we do know. For then, there would be an instance of knowledge for which, as we saw, (T)' doesn't hold.

More generally, if the inclusion of (T)' in the analysis of knowledge has the consequence that we wouldn't know many of the singular propositions we do know by induction, then (T)' has to be given up. Otherwise, Nozick is securing the failure of the Closure Principle, and, hence, freedom from Cartesian skepticism, at the cost of a strong kind of skepticism about induction. That hardly seems acceptable. It may be, though, that Nozick sees no *general* incompatibility between his tracking condition and inductive knowledge about particulars. He may think that conflicts arise only for a restricted, suspect class of propositions, one member of which would be 'The earth will not stop spinning suddenly.'¹⁷ At any rate, I propose to give Nozick the benefit

• PART
rt No-
NOWL-
idence
idence
as D.
rge S.
A. Fu-
ig the
cking,
athan
SKEP-
Steven
icism,
Peter
en F.
icism,
st the
ces
istant
ersity,
terna-
com-

of the doubt about his case, and to see how his account handles other cases of inductive knowledge about particulars.

Let's turn to a fresh example. Imagine it's a hot day in August, say 95 degrees in the shade. Several hours ago, you left some ice cubes in a glass out in the direct sun, and since that time you've gone inside to get out of the heat. You think about the ice cubes, and it occurs to you that the ice you left outside must have melted by now. Despite the fact that you are not, at that moment, perceiving the shallow layer of water at the bottom of the glass, you know that the ice has melted.

If (T) is necessary for knowledge, then in this situation it has to be true that

- (4) If the ice cubes hadn't melted, you wouldn't believe, on the basis of your actual evidence, that they had.

Now, I think, on a natural assessment, (4) comes out false. If the ice cubes hadn't melted, you would have been sitting inside thinking that they had, on the basis of all your past experience with ice, heat, and the like. The impression that (4) is false is bolstered by the fact that (5) seems true:

- (5) If the ice cubes hadn't melted, you would have been very surprised to learn of their state.

Presumably, the reason (5) is true is that, even if the ice hadn't melted, you would still have your normal expectation that it had. If that is so, (4) will be false.

The intuitive grounds for holding (4) false in this situation are backed up by semantic theory. According to David Lewis's theory of counterfactuals, (4) is evaluated by considering the most similar world to the actual one where the ice cubes didn't melt. Lewis's theory further indicates that the most similar world is one which is identical to the actual world up until the time when the ice cubes actually begin to melt. Since the histories of the two worlds are identical up to that point, your evidence about ice cubes and your beliefs formed from that evidence will be the same in both worlds. Thus, in the nearest possible world in which the ice cubes don't melt, you do form the belief that they have melted, on the basis of your actual evidence. On Lewis's theory, then, (4) is false for the example given.

Nozick is conscious of the incompatibility of his account of inductive knowledge with Lewis's semantics, and he has a response. In order to discuss it, I have to introduce some terminology. A *forward conditional* is one in which the antecedent refers to a time that is earlier than, or contemporaneous with, the time described in the consequent. So, 'If Tom had gone to Denver yesterday, he would have been snowbound today' is a forward conditional. A *backward conditional* is one which

says how counterfa
time later
Yale Coll
have beer
Finally, a
the comb
the forwa
Open fina
him.' Tha
but, view
come out
McEnroe
other, ver
forward
earlier ro
points fro
gether ma
nals, I w
condition

Lewis's
of condit
he needs
makes it
factuals i

The re
determ
identic
not-p.
such th
modifi

Nozick r
might lo
Still, it
ample, th
compou

(4) If
th

is suppo

(6) If
pr

says how things would have to have been before hand for some counterfactual situation to have come about; the antecedent refers to a time later than the time described in the consequent. An example is 'If Yale College had graduated 2,500 seniors in 1985, it would have to have been the case that they admitted over 2,000 freshmen in 1980.' Finally, a *backtracking compound* is a forward conditional treated as the combination of a backward and a forward conditional. Consider the forward conditional 'If I had played John McEnroe at the U.S. Open finals, I would have won a respectable number of points from him.' That may seem unlikely, in view of my lamentable tennis game, but, viewed as a backtracking compound, this counterfactual can come out true. For, we have the backward conditional 'If I had played McEnroe at the finals of the U.S. Open, I would have to have beaten other, very good players in the earlier rounds.' This combines with the forward conditional 'If I had beaten other, very good players in the earlier rounds of the Open, I would have won a respectable number of points from McEnroe in the finals.' Putting the two conditionals together makes it come out true that 'If I had played McEnroe in the finals, I would have won a respectable number of points,' where this conditional is taken as a backtracking compound.

Lewis's semantics excludes this kind of backtracking interpretation of conditionals, but such an interpretation is just what Nozick thinks he needs for his account of inductive knowledge. In a footnote, Nozick makes it clear that he wishes to reject Lewis's account of counterfactuals in favor of one that admits backtracking compounds:

The relevant not-*p* world [that is, the one to be considered when determining whether the tracking conditional is satisfied] is not a world identical to the actual one until now, and then diverging so as to produce not-*p*. . . a more complicated theory will be needed; my hope is that any such theory can be plugged into the text above, producing only minor modifications in it. (PE, 223n).

Nozick makes some suggestive remarks about what a suitable theory might look like, but these give no guidance in treating concrete cases.

Still, it is clear enough what Nozick has in mind. In the ice cube example, the relevant instance of (T) is to be treated as a backtracking compound. That is to say

- (4) If the ice cubes hadn't melted, you wouldn't have believed, on the basis of your actual evidence, that they had.

is supposed to be true because a backward conditional like (6) holds:

- (6) If the ice hadn't melted, it would have been the case that your previous experience was different in ways that would have led

you not to expect that the ice would melt under the circumstances.

My point is this: (6) is not clearly true at all. Who knows that your past would have been like if the ice cubes hadn't melted? Perhaps, if there were this one exception to the usual laws about the way ice behaves in the heat, there would have been many others. These might, then, have come to your attention beforehand, and led you not to expect the ice cubes to melt in the sun. On the other hand, it could just as well have been that the failure of the ice to melt on this occasion was a very rare, even unique, exception to the ways ice normally behaves. In that event, you would have had the same basis for expecting the ice to melt as you actually did. (6) is, at best, indeterminate, because there is no reason to favor one story rather than the other. Now, if (6) is indeterminate in this way, and (4) is true only if (6) is, then (4) is itself indeterminate, if it isn't false.¹⁸ Because (4) is at best indeterminate, (T)' would not be satisfied in this case, despite the fact that you do know by induction that the ice cubes have melted. (T)', then, cannot be necessary for knowledge.

The problem here is a general one, and it has to do with the role Nozick assigns to backtracking compounds. The backward conditionals they introduce are as a rule indeterminate.¹⁹ As a result, a high degree of indeterminacy will characterize the relevant instances of (T)' construed as backtracking compounds. So, if Nozick's approach to the evaluation of the conditionals were correct, and satisfaction of (T)' were necessary for knowledge, then attributions of inductive knowledge would themselves almost always be indeterminate (or false). Since such indeterminacy doesn't exist, Nozick has to be wrong about the status of (T)'.

Let me summarize the criticisms I have offered. If instances of (T)' are not taken as backtracking compounds, then (T)' fails outright in cases of inductive knowledge like the ice cube example. If, to avoid this result, the crucial conditionals are taken as compounds, attributions of inductive knowledge turn out to be, if not false, then highly indeterminate. The end result is that an account of knowledge which makes (T)' a necessary condition will not adequately handle knowledge by induction.

In constructing these objections, I have been relying, in the first place, on our intuitions about knowledge and about the truth-values of various conditionals. My point can be put this way: an inconsistency results from combining Nozick's tracking requirement, our intuitions about counterfactuals, and our intuitions about inductive knowledge. Under the circumstances, the tracking requirement ought to be given up. Sometimes, it appears, it is philosophically desirable to sacrifice intuitions in favor of some theory. I see no reason why such a step would

Tracking,

be justified
factually
would have
counterfactual
done if

VI. KN

My criticism
consider
position
stance, I
that uno
want to
point of
tics do I
and per
tracking

Here
the situa
been a v
degrees
basic ph
main fro
laws in t
tual one
hand, in
many o
differen
evant to
afterno
true in

(6) I
F
y
s

Moreo
after th
backwa
the ice
while,
force.²
happen

be justified here. Moreover, rejecting the intuitions about counterfactuals in particular in order to save Nozick's epistemological theory would have a theoretical cost as well. For insofar as Lewis's theory of counterfactuals matches those intuitions, it too would have to be abandoned if the tracking condition is retained.²⁰

VI. KNOWLEDGE AND NATURAL LAWS

My criticisms of Nozick in the previous section have been based on considerations concerning inductive knowledge of singular propositions. More specifically, I am claiming that, in a representative instance, Nozick's tracking condition is inconsistent with our knowing that unobserved exceptions to natural law don't occur. In this section, I want to take up an extended objection to the argument I just gave. The point of the objection is that, while some intuitions and Lewis's semantics do tell against the tracking condition, a revised semantic theory and perhaps some counterintuitions may be brought to support the tracking requirement.

Here is the objection: Let us reconsider the ice cube example. If, in the situation described, the ice cubes hadn't melted, there would have been a violation of the law that water melts at temperatures above 32 degrees F. The melting point of water, though, is determined by more basic physical laws. So, in the possible world where the ice cubes remain frozen, the normal physical laws could not have been in force; the laws in force in that world would have to be quite different from the actual ones. This difference, it might seem, would have shown up beforehand, in that world, in the previous abnormal behavior of water and many other things. In that event, your past experience would have been different from your actual experience in ways that would have been relevant to your belief about the state of the ice cubes on that particular afternoon. That is to say, there are reasons to think that (6) is probably true in this case:

- (6) If the ice hadn't melted, it would have been the case that your previous experience was different in ways that would have led you not to expect that the ice would melt under the circumstances in question.

Moreover, if we are concerned with your belief about the ice cubes *after* the fact, as you sit inside without having looked, the appeal to backward conditionals like (6) can be dispensed with altogether. For, if the ice cubes hadn't melted after having been out in the sun for a long while, at that point at least, the normal physical laws would not be in force.²¹ If they weren't in force, all kinds of strange things would be happening that you would surely notice. So, it is reasonable to assume

y • PART
ert No-
KNOWL-
vidence
vidence
mas D.
orge S.
A. Fu-
ng the
acking,
nathan
SKEP-
Steven
ticism,
Peter
hen F.
ticism,
ist the
xes
sistant
ersity,
terna-
hcom-

that, in the face of all sorts of apparently impossible phenomena, you would not then be inclined to hold the belief that the ice cubes were melted.²² This last counterfactual is not a *backtracking* compound; it combines the forward conditionals

- (7) If the ice weren't melted at time *t*, then at time *t*, the normal laws of nature wouldn't be in effect.

and

- (8) If, at time *t*, the normal laws of nature weren't in effect, you would be struck by the abnormal course of macro-level phenomena and not expect the ice to melt.

to yield the permissible forward conditional

- (9) If the ice cubes weren't melted at time *t*, you wouldn't expect the ice to behave as it normally does; that is, you wouldn't believe that it was melted.

On the face of it, though, this reply accomplishes little. It presupposes that if there were a local exception to the melting-point law, there would have to be widespread and noticeable exceptions to the more basic laws which underlie the law about the melting point. Why wouldn't the basic laws be broken only as they affect the ice cube at that time, leaving virtually everything—including your expectations—unaltered?

The proponent of (T') might answer as follows. One feature of the actual world is that it is governed by exceptionless physical laws that are invariant over space and time.²³ On this basis, it could be claimed that a possible world which has different, but unbroken and uniform physical laws more closely resembles the actual world than does a world governed by laws that admit exceptions. If similarity in this particular respect is given so much importance that it outweighs all other considerations, then the possible world(s) which determine the truth-values of the counterfactuals which concern us will be worlds governed by exceptionless laws. Specifically, the closest possible world to the actual one in which the ice cubes don't melt will turn out to be one where the basic physical laws permit ice to remain frozen at temperatures like 95 degrees F. Since these laws must be invariant over time and space, it is very plausible to think that your experience in that world would have been different enough, either before or at the time in question, to have led you not to expect that the ice left outside would melt. The key counterfactual 'If the ice hadn't melted, you wouldn't have believed that it had' comes out true, and (T') is satisfied here as Nozick requires.

Clearly, a working similarity ordering suitable for evaluating instances of (T') is not fully defined by the bare stipulation that possible

worlds
world
to con:
favor
worlds
plete c
all oth
hold.²
and a
the la
count
accou:
true:

(10)

The p
Des
for a
poner
">>
vant i
actua
as a s
cordi
holds
der th
No
fort v

(11

wher
wise.
that
can s
tions

(12

(It is
to th
alds
for (

worlds governed by uniform, exceptionless laws are closer to the actual world than worlds without such laws. It appears, then, that we ought to construe the argument of the previous paragraph as an argument in favor of modifying the standard similarity ordering for possible worlds given by Lewis. The modification consists in weighting complete conformity to invariant natural laws so heavily that it dominates all other criteria of similarity; otherwise, the criteria given by Lewis hold.²⁴ There will be substantial differences between Lewis's semantics and a semantics for counterfactuals built on the modified ordering. If the latter is taken as a *replacement* for Lewis's general theory of counterfactuals, an inferior semantics results. For example, on Lewis's account the following, apparently true counterfactual does come out true:

- (10) If I changed dolphins into men, I would have performed a miracle.

The modified semantics makes this come out false.²⁵

Despite the fact that the revised similarity ordering is inappropriate for a general semantics of counterfactuals, it remains open to the proponent of (T)' to use that semantics to define a new modal connective " $> > >$." On this proposal, when S knows that p , S satisfies the relevant instance of $[not-p > > > not-(S \text{ believes that } p \text{ on the basis of } S\text{'s actual evidence})]$. Here, though, the resulting sentence is *not* to be read as a standard English counterfactual. Rather, it is to be evaluated according to the specifications given: it is true just in case the consequent holds in the nearest $not-p$ world(s), where "nearest" means "nearest under the revised similarity ordering."

Now, even if we are willing to go this far in order to save (T)', the effort will be unavailing. Consider

- (11) The next emerald you see will not be "grue"

where "grue" means "green before the year 2,000 and blue otherwise."²⁶ Unless we succumb to skepticism about induction, we will say that (11) is something a person may know now. However, no subject can satisfy (T)', even in its most recent transmogrification, for propositions like (11). (T)' requires

- (12) If the next emerald you see were "grue," you would not believe, on the basis of your actual evidence, that the next emerald you see will not be "grue."

(It is understood that this counterfactual is to be interpreted according to the revised nonstandard semantics). The existence of "grue" emeralds is inconsistent with actual physical law. By the revised semantics for (T)', then, we must evaluate (12) by considering the nearest possible

y * PART
bert No-
KNOWL-
vidence
vidence
mas D.
orge S.
A. Fu-
ing the
acking,
nathan
SKEP-
Steven
ticism,
Peter
hen F.
ticism,
ist the
xes
sistant
ersity,
terna-
hcom-

world with "grue" emeralds that is also governed by exceptionless, time-invariant physical laws — if such a world exists. Otherwise, we are to follow the other criteria Lewis gives for evaluating counterfactuals.

Now, any possible world in which "grue" emeralds are found is *not* a world in which the laws of nature are spatio-temporally invariant. Hence, we must fall back on Lewis's semantics in considering (12). The nearest world, on Lewis's rendering, in which the next emerald you see is "grue" is a possible world exactly like the actual one, except that the emerald you will encounter has a propensity to change from green to blue in a little more than a decade from now. In this possible world, there is no reason why your experience beforehand should be any different from your experience in the actual world. So, you would have up until now all the same beliefs as you actually have, including those upon which you base your belief that the next emerald you see will *not* be "grue." It turns out, finally, that if the next emerald were "grue," you *would* believe, on the basis of your actual evidence, that it wasn't "grue." (12), then, is false, and (T)' fails for a genuine case of inductive knowledge.²⁷

If this example seems too exotic, consider an alternative. Suppose two policemen confront a mugger, who is standing some distance away with a drawn gun. One of the officers, a rookie, attempts to disarm the mugger by shooting a bullet down the barrel of the mugger's gun. (I assume that the chances of doing this are virtually nil). Imagine that the rookie's veteran partner knows what the rookie is trying to do. The veteran sees him fire, but is screened from seeing the result. Aware that his partner is trying something that is all but impossible, the veteran thinks (correctly as it turns out)

(13) The rookie missed.

Now, I would say, in this sort of situation, the veteran might well know that the rookie had missed. But (T)' requires

(14) If the rookie hadn't missed, the veteran would not have believed (by his actual evidence) that the rookie had missed.

where (14) is to be evaluated by considering the closest world, with invariant physical laws, in which the rookie does shoot a bullet down the barrel of the mugger's gun. It is possible, given the *actual* natural laws, for the rookie to succeed, although the likelihood of this is just about negligible. So, we may assume that the closest law-invariant world in which the rookie succeeds is one in which the actual laws hold; the invariance of those laws will do nothing to affect the veteran's reasonable expectation that the rookie will miss. As a result, in this world, the veteran believes the rookie will miss, and the consequent of (14) is false. However, as specified, the consequent of (14) is true, so the condi-

tional as a
vised con:

VII. CO

In the pr
Nozick's
tracking c
gued that
cases of ir
said on t
Principle
that princ

BIBLIOG

- Barker, S.
 ophic
Bennet, J.
 Revie
Garrett, B.
Goldman,
 (1983)
Goodman,
 Unive
Lewis, Da
 455-7
 . Ci
Luper-Fo
 Track
 (1984)
McGinn, C
 vol. C
 Minn
Nozick, R
 1981)
Shope, R
 to No
Stalnaker,
 Quar

NOTES

1. The
logical con:
might know
q because

tional as a whole comes out false. It emerges that (T)', even under its revised construal, is not satisfied in a case where someone would know.

VII. CONCLUSION

In the previous sections, I have been concerned to make clear how Nozick's denial of the Closure Principle rests on the claim that his tracking condition, (T)', is a necessary condition for knowledge. I argued that (T)' cannot be necessary for knowledge, because there are cases of inductive knowledge for which (T)' is not satisfied. If what was said on these points was correct, Nozick's case against the Closure Principle is unsuccessful. The skeptical argument which depends on that principle remains unrefuted.²⁸

BIBLIOGRAPHY

- Barker, S. F., and Peter Achinstein. "On the New Riddle of Induction," *Philosophical Review* 69 (1960): 511-22.
- Bennet, Jonathan. "Counterfactuals and Temporal Direction," *Philosophical Review*, 93 (1984): 57-91.
- Garrett, B. J. "Nozick on Knowledge," *Analysis* 43 (1983): 181-184.
- Goldman, Alvin. Review of *Philosophical Explanations*, *Philosophical Review* 92 (1983): 81-88.
- Goodman, Nelson. *Fact, Fiction, and Forecast*, 4th ed. Cambridge: Harvard University Press, 1983.
- Lewis, David. "Counterfactual Dependence and Time's Arrow," *Nous* 13 (1979): 455-75.
- . *Counterfactuals*. Oxford: Blackwell, 1973.
- Luper-Foy, Steven. "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism," *The Australasian Journal of Philosophy* 62 (1984): 26-49.
- McGinn, Colin. "The Concept of Knowledge," *Midwest Studies in Philosophy*, vol. 6, ed., French, Uehling, and Wettstein (Minneapolis: University of Minnesota Press, 1984).
- Nozick, Robert. *Philosophical Explanations* (Boston: Harvard University Press, 1981).
- Shope, Robert. "Cognitive Abilities, Conditionals, and Knowledge: A Response to Nozick," *The Journal of Philosophy* 81 (1984): 29-48.
- Stalnaker, Robert. "A Theory of Counterfactuals," *American Philosophical Quarterly*, monograph no. 2 (1968): 98-112.

NOTES

1. The exact formulation of the Closure Principle is complicated by various psychological considerations that have to be taken into account. For example, it is conceivable that S might know some proposition *p* and also know that *p* entails *q*, yet fail to draw the conclusion *q* because S fails to "put what he or she knows together." The formulation given doesn't

• PART
ert No-
NOWL-
vidence
vidence
nas D.
rge S.
A. Fu-
ng the
acking,
athan
SKEP-
Steven
icism,
Peter
hen F.
icism,
st the
kes
sistant
ersity,
terna-
hcom-

cover all these complexities. Problems of this sort, however, don't affect issues connected with skepticism, where the logical relations among the relevant propositions may be taken to be fully explicit to the subject. In such cases, where there is no question about whether the Closure Principle would apply, I call the entailed proposition a *clear* logical consequence of the proposition known.

2. Here and throughout I will use the symbol ' \rightarrow ' to stand for the subjunctive conditional; I use ' \Rightarrow ' to represent logical entailment. See David Lewis, *Counterfactuals*, and Robert Stalnaker, "A Theory of Conditionals" for a detailed discussion of subjunctive conditionals. Although Lewis and Stalnaker share a general approach, their analyses differ in important ways.

3. Nozick suggests that his theory may properly be viewed as a relevant alternatives account of knowledge. The application of the Lewis-Stalnaker framework to (A) and (T) provides a solution to the problem which plagues other relevant alternatives theories, viz. the problem of specifying which alternatives *are* relevant. For Nozick, the epistemically relevant situations, the ones in which S must be right about *p* in order to know that *p*, will be just those situations which must be considered in evaluating the subjunctive conditionals (A) and (T).

4. Alvin Goldman, *Review of Philosophical Explanations*, 84.

5. This formulation departs from the one given by Nozick (179), and incorporates certain refinements suggested by Steven Luper-Foy. Nozick's version of (T)' requires that S not believe that *p* in the nearest possible world where *p* is false *and* S does employ method M. By contrast, Luper-Foy's rendering makes a requirement concerning the nearest possible world where *p* happens to be false; this world need not be one in which S employs method M. However, if S does employ method M in that world, Luper-Foy's version of (T)' requires that S not believe that *p* by that method. See Steven Luper-Foy, "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Scepticism," p. 28-29.

6. See particularly the essays by Shope, Luper-Foy, Garrett, and the review by Goldman.

7. Since, in the discussion that follows, particular inference patterns don't play an important role, I generally suppress the part of the condition which deals with them.

8. It isn't difficult to verify that Nozick's belief that *J* satisfies all the other conditions of his analysis, not just the tracking condition. Therefore, Nozick not only *can* know that *J*, but does know that *J*, despite not knowing that not-SK.

9. The basic point here is attributed to Klein by Luper-Foy, op. cit., 45n. The example is Luper-Foy's.

10. A more plausible proposal along the lines suggested might be that S knows that *p* just in case (1) S tracks *p*; and (2) S tracks all of *p*'s *known* logical consequences. However, knowledge so defined fails to be closed under known logical implication. To see this, suppose the following: (1) S tracks *a*. (2) S knows that ($a \Rightarrow b$). (3) S tracks *b*. (4) S knows that ($b \Rightarrow c$). (5) S doesn't track *c*. Because of (5), S can't know that *b*. But S may not realize that *c* is a logical consequence of *a* as well as *b*; in that case, (5) doesn't impair S's knowing that *a*. So S knows that *a*, but doesn't know its clear logical consequence *b*.

11. Shope, "Cognitive Abilities, Conditionals, and Knowledge," 41.

12. Shope, 41n.

13. Among S's false beliefs would have to be all those that entail that he is someone with normally functioning senses negotiating his environment in the usual manner. It is logically possible that, by a coincidence of cosmic proportions, S's other forced beliefs about the world might happen to be true. Thus, Shope's hypothesis, as modified, seems to be a somewhat weaker version of the usual brain-in-the-vat scenario.

14. Shope cites some passages which, he thinks, show that Nozick believes we do know that (C), and presumably, I would add, (C)' as well; see Shope, 40. Unfortunately, Shope's citation is garbled, so his reading can't be verified. But, at any rate, if Nozick held the view Shope mentions, Shope's example shows that it was ill-considered.

15. Nozick might be expected to trace any resistance to this response back to devotion to the Closure Principle itself. Otherwise, what else could be counterintuitive about it? However, rejecting (T) as a necessary condition for knowledge because it conflicts with the Closure Principle is circular, if the rejection of (T) is supposed to be part of an argued defense of that principle.

16. Colin McGinn has offered a counterexample which, if acceptable, would show that tracking isn't necessary for knowledge. But as McGinn himself notes, intuitions about the

example c.
See McGinn
advanced
these have

17. No

"cooked-u

18. Th

that you

have belie

they were

19. Le

about ho

"Counter

20. Jo

counterfa

"Counter

against L

Nozick's

21. Th

ceased to

state of th

I treat the

22. Si

in this w

elsewhere

23. Th

think tha

24. Se

25. Th

cubes hac

treatment

26. Se

given by

27. It

that emer

from acti

comes ou

28. I

Anthony

Mandel,

example conflict, so his criticism of the tracking requirement is, at best, not fully conclusive. See McGinn, "The Concept of Knowledge," 531-2. Also, I understand that Saul Kripke has advanced arguments against taking tracking as a necessary condition for knowledge, but these have not yet appeared in print.

17. Nozick might insist that 'The earth will stop spinning *suddenly*' is in some sense a "cooked-up" or skeptical hypothesis, the denial of which we really don't know.

18. The indeterminacy of (6) might make (4) outright false, if that indeterminacy means that you *might* have believed the ice cubes had melted when they hadn't. For if you might have believed that they were melted, it is false to say that you would not have believed that they were melted, in the situation described.

19. Lewis has observed that "Seldom, if ever can we find a clearly true counterfactual about how the past would be different if the present were somehow different." Lewis, "Counterfactual Dependence and Time's Arrow," 455.

20. Jonathan Bennett and others have criticized both Lewis's general approach to counterfactuals and his views about backward conditionals in particular. See Bennett, "Counterfactuals and Temporal Direction." I am not convinced by the objections raised against Lewis; also, it is my impression that the alternative accounts do not vindicate Nozick's epistemology.

21. There is a more complicated possibility, namely that the laws might have changed or ceased to be in effect temporarily, but returned to normal shortly before you reflected on the state of the ice cubes. This more complicated possibility could still be treated in the same way I treat the simpler situation in the text.

22. Since transitivity fails for counterfactuals, it isn't always permissible to combine them in this way. However, as I go on to claim, the real problem with this line of thought lies elsewhere.

23. This may not, as a matter of fact, be completely accurate since there is some reason to think that the values of basic physical constants change over time.

24. See Lewis, "Counterfactual Dependence and Time's Arrow," especially 472.

25. There is also a discrepancy between the two accounts as to the truth of (5) 'If the ice cubes hadn't melted, you would have been surprised by their state.' My sense is that Lewis's treatment, which makes (5) true, is correct.

26. See Nelson Goodman, *Fact, Fiction, and Forecast*, 74. I follow the interpretation given by Barker and Achinstein in "On the New Riddle of Induction." See 511.

27. It might be thought that there could be a world where it was a law, true at all times, that emeralds are "grue." In 1986, such a world would be indistinguishable to its inhabitants from actual world, and they would believe as we do that emeralds are not "grue." So, (12) comes out false anyway.

28. I wish to thank a number of people for helpful conversations on earlier drafts: Anthony Brueckner, John Martin Fischer, Harry Frankfurt, Steven Luper-Foy, Ross Mandel, and David Shatz.

THE POSSIBILITY OF KNOWLEDGE

NOZICK AND HIS CRITICS

Edited by

Steven Luper-Foy

Trinity University

ROWMAN & LITTLEFIELD
PUBLISHERS

1987