Nozickian Epistemology and the Question of Closure

A distinctive aspect of Nozick's contribution to the epistemology of the last half of the twentieth century concerns the question of whether knowledge is closed under known implication. Here I will examine the implications of this approach to the nature of knowledge regarding the question of closure. Nozick notoriously reveled in the failure of closure on his account, raising the question whether such merriment is warranted. Moreover, some have claimed recently that minor emendations in Nozick's approach will save closure. I will argue that the question of closure provides a serious obstacle to Nozickian approaches to epistemology.

Nozick's theory of knowledge clarifies the conditions for knowledge beyond true belief by appeal to counterfactuals, and there are four counterfactuals on which a theory might focus in giving a counterfactual-based account of knowledge. Where 'p' is a sentence, 'B' is the operator 'S believes that', and '$\Box\!\!\rightarrow$' the symbol for counterfactual implication, the four candidates are, with identifying descriptions:

1. $p \Box\!\!\rightarrow Bp$ (the receptivity condition)
2. $Bp \Box\!\!\rightarrow p$ (the safety condition)
3. $\sim p \Box\!\!\rightarrow \sim Bp$ (the sensitivity condition)
4. $\sim Bp \Box\!\!\rightarrow \sim p$ (the idealism condition)

Neither the idealism condition nor the receptivity condition have been thought to have much plausibility as conditions for knowledge in recent epistemology, and for good reason. The receptivity condition would limit what we know to things that wouldn't be true unawares. Much of what we know requires attention on our part, however. The fact that don't know what is going

on around us when we are asleep doesn't imply that we don't know what is going on around us when we are awake and attending to the features of our immediate environment. Moreover, even when we are alert, we are sometimes aware of details of our environment that we need not have noticed–the butterfly on the flower below my window, for example. In such cases, the claim could easily have gone unnoticed, and thus is a truth to which I am not receptive. Since I did notice the butterfly in this case, however, I do know that there is a butterfly on the flower below my window.

A similar problem affects the idealism condition. We can know what is going on around us when we are awake and attending to the features of our immediate environment, and yet the idealism condition limits what we can know to what would be false when not believed. That would require that anything about our immediate environment that is unchanged when we are asleep would fail to be a candidate for knowledge. In addition, it would require that anything about our environment that would be unchanged if unnoticed would also fail to be a candidate for knowledge. Just as with the receptivity condition, this condition is far too strong a condition on what we know.

Such considerations explain the focus of recent epistemology on the other two conditionals, the safety and sensitivity conditions. Nozick relies primarily on the sensitivity condition, while epistemologists such as Ernest Sosa and Timothy Williamson prefer the safety condition.[1] Both approaches are Nozickian in the sense that they distinguish knowledge from true belief on the basis of subjunctive conditionals. I will begin by comparing these two approaches on the question of closure.

Closure

A theory of knowledge respects closure when it adopts the idea, in some form, that the logical consequences of what we know are also known. This rough closure claim is certainly false, otherwise we'd all be logically and mathematically omniscient if radical skepticism is false. Theories that wish to respect closure defend the view by restricting this rough form of the view. Such restrictions can proceed either by weakening the antecedent–say, by requiring only that one be in a position to know, rather than that one actually knows–or by strengthening the antecedent–say, by requiring that one know that what one knows entails some further claim, rather than the entailment claim merely being true.

The project of refining the closure claim so that it is defensible is not a simple and easy project, but there is intuitive plausibility to the general idea, for we often appeal to some kind of closure condition. One example is in explaining why a claim to knowledge is mistaken. When identifying a particular SUV as a Ford Explorer, my assertion was challenged by one of my children with the question, "How do you know it's not a Nissan?" The point of the question is that Nissan has an SUV with the same body style as the Ford Explorer, and the presupposition of the question is that I could only know that it was an Explorer if I knew it was not a Nissan. This presupposition appears to involve some sort of closure principle about knowledge, one that I accepted in the context, because I replied that I saw the Ford logo on the front hood. We also often appear to extend our knowledge by applying deductive principles, thereby requiring the truth of some closure principle or other. Closure is thus initially attractive precisely because it provides a convenient explanation of examples of this sort.

Nozick noticed and embraced the fact that the sensitivity requirement does not preserve closure with respect to knowledge, using this feature of his account to respond to skepticism. Even if we take the simplest one-step derivations, Nozick theory rejects closure. For example, it is hard to think of a simpler one-step derivation than &-Elimination (&E), according to which we are allowed to infer *p* from *p&q*. Nozick rejects closure for even such simple derivations:

> So, we must adjust to the fact that sometimes we will know conjunctions without knowing each of the conjuncts. Indeed, we already have adjusted. Let *p* be the statement that I am in Emerson Hall, not-SK be the one that I am not on Alpha Centauri floating in that tank; since *p* entails not-SK, *p* is (necessarily) equivalent to *p* & not-SK. I know that *p*, yet I do not know that not-SK.[2]

There are two points to note here. First Nozick draws a conclusion based on the details of his theory that simple one-step derivations such as those involved in the application of &E do not preserve closure. The second point is more intuitive, and thus independent of Nozick's theory to a certain degree. Nozick says in the quote above, "Indeed, we already have adjusted," and then gives an account of why knowledge is not closed under applications of &E. The discussion of the case of being in Emerson Hall and not being on Alpha Centauri floating in a tank is intended to show that there are intuitive grounds for rejecting such a closure principle. Nozick thus holds both that his theory undermines closure, but also that it was undermined already, on intuitive grounds.

Nozick's particular example is not as compelling as one could wish. The most common reaction by ordinary folk to far-fetched skeptical scenarios is simply to deny that such scenarios undermine knowledge. Those more skeptically inclined may challenge this knowledge claim,

and an adequate answer to the challenges may be hard to find, but the absence of an adequate response to skeptical challenges does not undermine knowledge.

There are other examples, however, that are more persuasive. Suppose I've been invited to give a talk in St. Louis next month and have agreed to go, and my brother calls to ask if he can visit the very same date as the talk. I reply that there is a problem since I won't be home that day. That I won't be home that day is something I know. Yet, I also know that tragedies sometimes happen–people have heart attacks, strokes, auto accidents, etc.; and I also know that were I to suffer some of these tragedies, I would be at home recuperating. But I don't know that I won't suffer these tragedies.

Lottery situations provide other examples. My daughter sees a million-dollar mansion on television and wants to know if we could buy a house like that. I reply that I'll never have enough money for such a house, and this claim is something I know to be true–after all, I'm a philosophy professor, happy in this role with no intentions of changing careers, and I have no rich relatives from whom I might receive a large inheritance. I bought a lottery ticket, however, and I know that if I win the lottery, I would be able to buy such a house. I may have sufficient evidence for thinking that I won't win the lottery, but I would not claim to know that I won't win the lottery.

It is important to note that not all such cases involve knowledge of the future. One of the original examples in the literature on this problem involves ordinary perception. You visit the zoo in San Diego with your children, and inform your children that what they are seeing is a zebra. You also know that if it is a zebra, it is not a cleverly disguised mule painted to be indistinguishable to the ordinary observer from a zebra. You do not know, however, that the

animal you are looking at is not a cleverly disguised mule.[3]

Cases such as these are cases that appear to threaten closure principles about knowledge, and a significant literature has built up over the past thirty years or so about such cases.[4] In each such cases, there is an *ordinary proposition*, the proposition to which we naturally ascribe knowledge, and an *implied proposition*, the proposition implied by the ordinary proposition and to which we are not ordinarily willing to ascribe knowledge.

It is worth noting, however, that there is a variety of approaches to such cases that attempt to accommodate the idea that we know, e.g., that we won't be able to afford a mansion and yet do not know that we will not win the lottery without abandoning closure. One such approach is Contextualism, according to which the truth value of attributions of knowledge depends on the standards in place at the context of utterance. Some contexts are high standards contexts and some are ordinary standards contexts, leaving open the possibility that shifts in context can occur to explain how it could be true to say that one knows that one won't be able to afford a mansion and yet not to be true to say that one knows one will not win the lottery. The two most prominent defenders of Contextualism, Stewart Cohen and Keith DeRose, both maintain that one of the benefits of their view is that it preserves closure.[5]

Other views have a similar response. An alternative defended by John Hawthorne, which he terms a version of Invariantism, also allows the truth value of knowledge ascriptions to vary, where the variability traces to the situation of the individual whose knowledge is in question rather than to the situation of the individual ascribing knowledge, as it does for the Contextualist.[6] In general, pragmatic theories of knowledge, according to which the standards for knowledge are raised when the issue in question of truth is more important from a practical

point of view or more salient from the point of view of the cognizer in question, can explain apparent cases of failure of closure in terms of a shift in the standards for knowledge. Because the shift in standards that matters is that of the individual to whom knowledge is ascribed, such theories are Invariantist, since the central tenet of Contextualism is that the relevant shift involves the ascriber of knowledge rather than the individual to whom knowledge is ascribed.

By appealing to shifting standards in these ways, both views are able to avoid abandoning a closure principle about knowledge. An appropriate closure principle will be restricted to a situation in which the standards for the truth of knowledge ascriptions remains the same. So when shifting standards explain how an ordinary proposition can be known even though the related lottery proposition is not known, the shifting of standards prevents such cases from being counterexamples to an appropriate closure principle concerning knowledge.

It is worth noting that no fancy semantical theory is needed to offer a theory with these results. Both Contextualists and Invariantists emphasize various ways in which the standards for knowledge can be raised, one important item being the salience of the possibility of error.[7] If we think of the concept of knowledge as involving the concept of justification, details of the theory of justification can yield the suggested implication without any recourse to complex semantics posited by Contextualist (and some Invariantist) theories. One promising approach to the theory of justification treats justification as a function of positive evidence and known defeaters: justification amounts to adequate positive evidence not ultimately undermined by defeaters of which one is aware. In order to accommodate the Contextualist emphasis on salience of the possibility of error, one need only treat the defeater condition so that concerns by the individual in question about the likelihood and significance of error are themselves defeaters. That is, if the

possibility of error is salient for an individual regarding a particular claim they are considering, then that individual has a defeater of the positive evidence possessed for the claim in question. Such an account is compatible with a contextualist semantics for 'knows' but does not require it.

One perceived benefit of treating these standards semantically is the capacity to preserve a closure principle about knowledge. When the semantic factor regarding the standards for knowledge is held fixed, a closure principle can be maintained without denying that, e.g., we can know we have hands even though we do not know that we are not brains in a vat. Treating these factors as a special feature of the concept of defeat yields the same preservation of closure. The closure principle, on such a theory, will be restricted in terms of these special features of the concept of defeat, but only implicitly (since I'm assuming that justification is necessary for knowledge).

In sum, there are various approaches to these puzzles that allow for some closure principle regarding knowledge. So the point to note is that Nozick's appeal to intuition, by itself, cannot show that closure must be abandoned.

Moreover, the abandonment of closure is ordinarily viewed as a defect of Nozick's approach, and others have tried to maintain the spirit of a Nozickian approach to knowledge by replacing sensitivity conditions with safety conditions, hoping thereby to preserve closure.[8]

Consider the following two arguments:

1. p entails q	1. p entails q
2. ~p □→ ~Bp	2. Bp □→ p
3. So, ~q □→ ~Bq	3. So, Bq □→ q

The first column represents the argument that would need to be valid in order for sensitivity to preserve closure, and the second column the argument need for safety to preserve closure. We know that the first argument fails, on the basis of the Nozick's example about being in Emerson Hall and not being a brain-in-a-vat. A bit simpler example to follow concerns having hands and being in a Cartesian evil demon world. Both safety and sensitivity counterfactuals are true: if I had no hands, I wouldn't believe that I do; and if I were to believe that I have hands, I would have hands. My having hands entails that no demon is active who makes it always and everywhere false that anyone has hands. But if there was such a demon, I wouldn't believe that there was, so this belief is not sensitive to the facts even though it is entailed by one that is. Such an example does not constitute a counterexample to the safety requirement, however, for if one were to believe that there is no demon, there would be no demon, (given the operating assumption that the actual world is not a demon world), so this belief is safe as well. So the standard examples used to show that sensitivity is not closed under entailment lend support to the idea that safety is so closed.

This case highlights one of the features of the logic for counterfactuals. One of the truth-preserving operations on material conditionals is transposition (contraposition), but this rule is one of the standard rules that are invalid when used in counterfactual contexts.[9]

Still, concluding that safety is closed under entailment would be incautious. First, note that it would an amazing coincidence if every case in which sensitivity fails the closure test constitutes a counterexample to transposition on counterfactual conditionals, and yet that is precisely that to which defenders of closure for safety are committed. This intuitive point against the claim that safety involves closure can be further supported using the standard semantics of

possible worlds and similarity relations for counterfactuals, for this apparatus makes it very easy to construct a countermodel to both arguments above. Stipulate that in the actual world α the following is the case:

True-in-α: p, Bp, q, ~Bq, p entails q.

This stipulation guarantees the truth of the first premise of the above arguments. Further stipulate that in all worlds **W** most similar to α, we have:

True-in-W: ~p, ~Bp, ~q, Bq.

On this model, the premises of both arguments are true, and each conclusion is false. For the first argument, the conclusion is false since in the most similar world where q is false, Bq is true. For the second argument, the conclusion is false as well, since in the closest world in which Bq is true, q is false.

This result should not surprise us. We know that contraposition fails for counterfactual conditionals, but that fact is compatible with such contrapositives nearly always having the same truth value. So it shouldn't surprise us if the contrapositives defining safety and sensitivity have the same truth value in some special context, unless we had reason to think that such a context was one where counterexamples to contraposition were likely to occur. It is hard to find any basis for such a suspicion in the special context of closure issues, so we shouldn't expect safety to be immune from whatever closure failures bear on sensitivity.

One may still resist such a conclusion by pointing out general difficulties with the standard semantics for counterfactuals, and claiming that more intuitive readings of subjunctive conditionals on a case-by-case basis fail to produce counterexamples. The upshot of such a position is that even though we can't see some intuitive connection between closure difficulties

for sensitivity and counterexamples to contraposition, the case-by-case examination shows that there is nonetheless a connection between them. If these claims were true, that would give some reason for maintaining the closure claim about safety and blaming the inadequacy of the standard semantics for formal failures such as the one described above.

The problem is that the case-by-case claim made above is false. Suppose Bill is in fake barn country where all the fake barns are red, and the rare remaining real barn is green (imagine that this result is deliberate on the part of the inhabitants–they've decided to make every real barn be green and every fake barn be red). Then consider the claim that the building in the pasture is a green barn, and the related claim that the building in the pasture is a barn. Suppose that Bill is asked whether the object in question is a green barn, and Bill answers, "Yes," believing that the object in question is a green barn. His belief is safe: anytime he forms such a belief in an environment relevantly like the present one, he will be correct. Bill believes also that the object in question is a barn. This belief, however, is not safe: If he were to form such a belief in an environment relevantly like the present one, he would be mistaken.[10] Hence, safety does not preserve closure.

So safety theorists have no more right to claims about closure than do sensitivity theorists; the difference is that defenders of sensitivity have admitted failure of closure and safety theorists such as Sosa deny that their view has this implication. I think we can explain this difference in terms of a confusion between two closely related alternatives to sensitivity. I think defenders of closure using the language of safety have confused safety with a closely related view that is best expressed using the language of reliable indication. If one begins by thinking about sensitivity and safety in terms of subjunctive conditionals, one may be tempted to think of

reliable indication as expressible in terms of a subjunctive conditional where the reliable indicator is the belief itself. Thinking in this way generates the following closure argument, which I juxtapose with the failed safety closure argument:

1. p entails q	1. p entails q
2. Bp $\Box\!\!\rightarrow$ p	2. Bp $\Box\!\!\rightarrow$ p
3. So, Bp $\Box\!\!\rightarrow$ q	3. So, Bq $\Box\!\!\rightarrow$ q

Notice that this first argument, the argument representing closure of reliable indication, has no counterexamples, either intuitive or formal. For the second premise to be true, we must consider some situation in which both Bp and p are true, and the first premise requires that in any such situation q will be true as well. Thus, such a situation cannot be one where Bp is true while q is false.

Reliable indication, however, is a different condition from the safety condition, and defenses of closure under reliable indication do not support closure under safety. So no mere substitution of a safety condition for Nozick's sensitivity condition will save the view from the untoward failure of closure. Instead, what will be needed is a replacement of the kind of theory Nozick favors with a version of reliabilism. In doing so, one might attempt to highlight the resemblances between the two views by clarifying reliable indication in terms of counterfactuals, but one will be hard pressed to rely on the resemblances for long. For example, if believing that p is a reliable indicator of p, why think that it is the only kind of reliable indicator that can function in one's theory of knowledge? Why not any reliable indicator, whether a belief or not? For example, perhaps sensory experience itself is a reliable indicator of the truth of certain propositions. If so, one's theory should not rule out this kind of reliable indicator as suitable in

the account of knowledge?

Reliable indicator theorists may remain Nozickian by insisting that any reliable indicator of truth will make belief also be a reliable indicator of truth. Such a response generates a new violation of closure, for it commits the reliable indicator theorist to the following argument form:

1. p entails q
2. X $\Box\!\!\rightarrow$ p (assumption that X is a reliable indicator of p)
3. So, X $\Box\!\!\rightarrow$ q
4. So Bp $\Box\!\!\rightarrow$ p (using "if X is a reliable indicator of p, then so is believing p.")
5. So, Bq $\Box\!\!\rightarrow$ q (again using "if X is a reliable indicator of p, then so is believing p," with q as the value for p)

This argument implies a restricted form of closure for safety: it says that if you have a reliable indicator for p, then the belief that p is safe and so is any belief entailed by the belief that p. The barn example above shows that one cannot endorse such a view: even though (4) is true when p is about green barns, (5) is false when q is about barns. (1)-(4) may all be true and (5) still false.

So if an interest in preserving closure leads one to endorse a reliable indication theory, the best thing to do is simply embrace reliabilism and stop trying to pretend that one is defending a view of the same family as Nozick's. Restricting the reliable indicators to beliefs makes the theory insufficiently general, and once one generalizes, no basis is left for maintaining that the reliable indicator view is some small variant of the safety view. The conclusion to draw is that substituting a safety condition for Nozick's sensitivity condition provides no basis for rescuing Nozickian epistemology in the face of concerns about closure.

Probability and Closure

A natural move at this point is to wonder whether the closure of safety can be rescued by inserting talk of probabilities into these approaches. Instead of the above conditions, a theorist might say that a belief is safe when one would not easily be mistaken or when most of the time the claim in question is true in the presence of belief,[11] and one might say that a belief is sensitive when most of the time one would not hold the belief if the claim were false.

I assume that remarks about what one could easily be mistaken about are probabilistic in character, just as remarks about what is true most of the time are probabilistic in character. Interpreted in this way, such claims contain a scope amibiguity, depending on whether the probability operator is taken to govern the entire conditional or only the consequent of the conditional. Disambiguating gives us the following:

Wide scope probable safety: $Pr(Bp \ \Box\!\!\rightarrow p)$

Wide scope probable sensitivity: $Pr(\sim p \ \Box\!\!\rightarrow \sim Bp)$

Narrow scope probable safety: $Bp \ \Box\!\!\rightarrow Pr(p)$

Narrow scope probable sensitivity: $\sim p \ \Box\!\!\rightarrow Pr(\sim Bp)$

The first thing to notice is that closure does no better on the narrow scope probabilistic readings than on the non-probabilistic versions of safety and sensitivity. If Bill were to believe that the object in question is a barn, in all likelihood he'd be wrong–after all, the real barns in the area are extremely rare. So if our goal is to preserve closure, neither the move from sensitivity to safety, nor the move from narrow scope probabilistic sensitivity to narrow scope probabilistic safety will help.

What about wide scope readings? Such readings generate the following arguments:

1. p entails q	1. p entails q
2. Pr(~p □→ ~Bp)	2. Pr(Bp □→ p)
3. So, Pr(~q □→ ~Bq)	3. So, Pr(Bq □→ q)

These arguments are harder to assess in large part because of the variety of interpretations proposed for the Pr operator, and the accompanying difficulties for each. For example, if we take a frequency approach to probability, these arguments will fail since it will be a contingent matter whether the claim within the scope of the probability operator in the conclusion is true most of the time when the premises are true. A similar point holds for other interpretations of probability, such as subjective theories, propensity theories, and nomological ones, namely that the truth of the conclusion in face of the truth of the premises will be contingent at best. This point undermines the validity of these arguments, for arguments are valid only when the premises necessitate the conclusion.

The only approach that might allow the validity of these arguments is also one of the more mysterious ones, the logical theory, which maintains that any true probability judgment is necessarily true. If the conclusions are true, the above arguments will be valid since any argument with a necessarily true conclusion is valid.

Even so, there is cause for concern here. Recall that the defender of safety is claiming superiority for that view over the sensitivity view on the basis of closure results. So even if the second argument above is valid, that will not be a reason to favor safety over sensitivity unless the first argument is invalid. How would a defender of safety show that the first argument is invalid (assuming the logical theory of probability)? The best that can be done, I think, is to

show that the information in the premises leaves the likelihood of the embedded sentence of the conclusion utterly inscrutable, i.e., they provide us with no information from which we can obviously conclude that the conclusion must be true. The same holds, however, for the second argument. So the challenge to the defender of safety is that it is not enough to maintain that the second argument is valid. Some difference must be found between it and the former argument if the superiority of safety is to be sustained.

Moreover, the move to any probabilistic reading raises a further problem for the idea that some safety condition explains the difference between knowledge and true belief. Consider how beliefs in lottery cases can be probabilistically safe. No adequate theory of knowledge can imply that we know that our ticket in the lottery will lose, no matter how large the lottery is, but if some probabilistic safety condition distinguishes knowledge from true belief, then such knowledge is possible. In the case of a relatively large lottery where one believes that one's ticket will lose, one would not easily be false, and thus one's belief partakes of probabilistic safety. The lesson is that probabilistic safety theorists will need some further condition to give an adequate account of the nature of knowledge.

We must judge, therefore, that the probabilistic move is a non-starter. There is no reason to think it will do better on the closure issue, and it raises additional problems not faced by non-probabilistic approaches.

Deeper Problems for Safety and Sensitivity Conditions

In light of these conclusions, it might seem wise to embrace the failure of closure, as

Nozick does, hoping that the theoretical virtues of a counterfactual approach to the nature of knowledge are impressive enough to outweigh the discomfort caused by abandoning closure. Such an approach would have to be developed carefully, for some failures of closure are worse than others. For example, it is an intolerable failure for a theory to imply that one can know that something is a green barn, but not that it is a barn. Such a position is truly deserving of DeRose's appellation, "the abominable conjunction."[12] My own intuitive response is not nearly as hostile to the suggestion that we can know we have hands even though we do not know that we are not brains-in-a-vat being stimulated to experience the world in precisely the way we are actually experiencing it. I am not saying that I find the latter conjunction unproblematic. I am saying, though, that I find it less problematic than the former conjunction, and it may be that there is a way to develop a safety or sensitivity approach to reduce the kinds of failures of closure to those that are the least problematic kind.

It is worth noting as well the difficulties encountered in recent attempts to formulate a defensible closure condition.[13] In light of the difficulties in formulating an acceptable closure condition, we should be suspicious of categorical rejections of a theory based only on closure considerations.

I will argue, however, that there are deeper problems with Nozickian approaches to the nature of knowledge. I will argue that neither sensitivity nor safety approaches are suited for explaining the nature of knowledge, but instead can explain only part of what we know. The limited roles I will argue for are these: sensitivity is more naturally suited to explaining perceptual knowledge, whereas safety is more naturally suited for explaining statistical or inductive knowledge, and neither is well-suited for explaining the other category.

Sosa argues against the sensitivity requirement and in favor of the safety requirement in just this way. He imagines a case where a person has dropped one's garbage down a chute in an apartment building, and that person believes that the garbage will soon be in the basement. The long track record of correlation between the two events makes the belief safe: if one believes, in this situation, that the garbage will soon be in the basement, it would be true. Notice, however, that the belief is not sensitive. If the garbage would not soon be in the basement, it likely would somehow have gotten caught in the chute, in which case one would still believe that the garbage will soon be in the basement.

What is instructive here is that Sosa's defense appeals to a case of inductive knowledge, and sensitivity is not well-suited to such contexts. It is equally instructive to notice that things go in the other direction in perceptual cases. Suppose one is color blind to the extreme: everything looks brown. Suppose also that in one's local environment, everything is a shade of brown. Then, when looking at a brown thing, one's belief is safe: if one believes it is brown, it would be brown. But one's belief is not sensitive, and its insensitivity is the important point here, for such a person does not know the color of the thing he is looking at even if the environment is so unbelievably friendly to him that he is always right.

In perceptual cases, our attention generally focuses on powers of discrimination, and here sensitivity is central. A substantial track record of truth may show that a belief is safe, but it won't show that it counts as knowledge when one lacks the discriminatory capacity to differentiate truth from error. Lacking such a capacity makes one's beliefs insensitive, and this feature dominates the safe character of the belief, for in such cases, knowledge is lacking.

If I am right, neither sensitivity nor safety are conditions for knowledge. Sensitivity is

more at home in the realm of perceptual knowledge and safety in the realm of inductive or statistical knowledge. Of course, it is open to a theorist to opt for one of the two, and simply deny that there is knowledge of the other sort. Such a route would be a very difficult one for a defender of the safety condition, for that would require holding that inductive knowledge is possible even though perceptual knowledge isn't. More likely is the other route: to hold that perceptual knowledge is possible, but statistical or inductive knowledge is not. I will not pursue either line here, except to point out what I take to be the severest counterexample to the idea that inductive knowledge is not possible. One of the deepest aspects of the human condition is coping with our own impending death, and among the most important items of self-knowledge that we have is that we will die. Only theory-clouded judgment can maintain that I can know what color the sky is but not that I am mortal.

The other option is to bifurcate one's theory of knowledge, maintaining that one of the two counterfactual conditions we've examined always explains the difference between knowledge and true belief. Not so, as is shown by a central example in the Gettier literature.[14] You see Tim steal a book from the library, and believe as a result that Tim stole the book. You report Tim to the police, and when they go to Tim's home to arrest him, Tim's mother swears his innocence. She says that Tim has a twin brother, Tom, who is a kleptomaniac, and that Tim is on a safari expedition in Africa. Where is Tom? She claims to have no idea, of course.

This case splits into two. In one case, the police have no past history with Tim's mother. Her claims are persuasive and she appears honest, so the police must check out her story before proceeding with their plans to find Tim and arrest him. In the other case, the police have a long history of acquaintance with Tim's mother, knowing that she has lied repeatedly to save Tim

from trouble. In fact, she has used this precise lie in the past, and the police have investigated the story and learned that Tim does not have a twin but that he is an only child. In such a case, they laugh off the tale and carry on with their plans to arrest Tim.

In the second case, both you and the police have knowledge about who stole the book that is unaffected by the mother's testimony, but not in the first case. In both cases, the mother's testimony is a defeater of your evidence that Tim stole the book, but it has differential significance in the two cases. In the first case, it undermines your knowledge, since it is a possibility that the police must check out. But in the second case, it doesn't undermine your knowledge.

If this particular case isn't convincing, imagine different contexts for the mother to say what she says. Imagine, for example, that the mother's testimony is part of a theater production in which she was acting. Or imagine her saying it in her sleep, as part of a dream, or saying it to convince herself: "He *has* a twin brother! He really does! Tim's in Africa, he *really* is!" And imagine any of these possibilities with no one around at all to hear what she says. Knowledge can't be so fragile that it disappears just because a normally competent speaker denies what we believe.

Notice that neither safety nor sensitivity can explain the difference between the two cases. In both cases, the mother is failing to tell the truth, and whatever is the proper account of the relevant difference in the two cases, that difference cannot be explained in terms of safety or sensitivity.

Furthermore, notice that there will be cases like this one for any type of belief one wishes to consider. The above case involves a perceptual belief, but it could just as easily have been a

statistical or inductive belief, a belief based on memory, or a testimonial belief. Safety and sensitivity may be part of the story of what distinguishes knowledge from true belief, but they are not enough of the story to fulfill the promise of Nozickian epistemology. Neither are necessary for knowledge, and together they are not sufficient either.

NOTES

1.. See Ernest Sosa, “Proper Functionalism and Virtue Epistemology,” and “Postscript to “Proper Function and Virtue Epistemology,”” both in Jonathan L. Kvanvig, ed., *Warrant in Contemporary Epistemology*, (Lanham, Maryland: Rowman & Littlefield, 1996), pp. 253-270 and 271-280; and Timothy Williamson, *Knowledge and Its Limits*, (Oxford: Oxford University Press, 2000).

2. Nozick, *op. cit.*, p. 228.

3.. See Fred Dretske, “Epistemic Operators,” “Conclusive Reasons,” *Australasian Journal of Philosophy* 49, (1971), pp. 1007-1023, and “The Pragmatic Dimension of Knowledge,” *Philosophical Studies* 40, (1981), pp. 363-378.

4.. See John Hawthorne, *Knowledge and Lotteries*, (Oxford: Oxford University Press, 2004) for a discussion of the problem and a thorough guide to the literature.

5.. See, e.g., Stewart Cohen, “How to be a Fallibilist,” *Philosophical Perspectives* 2 (1988), pp. 581-605; “Skepticism, Relevance, and Relativity,” in *Dretske and His Critics* (Oxford: Basil Blackwell, 1991), pp. 17-37; “Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery,” *Australasian Journal of Philosophy* 76, 2 (1998), pp. 289-306; Keith DeRose, “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52 (1992), pp. 913-929; and “Solving the Skeptical Problem,” *The Philosophical Review* 104, 1 (1995), pp. 1-52.

6.. Hawthorne, *Knowledge and Lotteries*.

7.. Cohen is most direct about the importance of such salience, and this feature of standards-

raising is appealed to in a way that generates, to my mind, disastrous consequences by David Lewis in "Elusive Knowledge," *Australasian Journal of Philosophy* 74, 4 (1996), pp. 549-567, some of which are noted in later discussion in the text. According to Lewis, merely reflecting on the question of whether we have met the requirements for knowledge causes the salience condition to raise the standards to the impossible heights, so that such reflection always and everywhere destroys the possibility of knowledge–more carefully, it raises the standards so that every statement of the form 'S knows that p' is false.

8.. A clear and helpful discussion of these issues can be found in Steven Luper, "The Epistemic Closure Principle", The Stanford Encyclopedia of Philosophy (Spring 2002 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2002/entries/closure-epistemic/>.

9.. See, e.g., David Lewis, *Counterfactuals*, Oxford: Basil Blackwell, 1973.

10.. I thank Stew Cohen for this point, who noted that the example is borrowed from what are now referred to as Saul Kripke's "Nozick-bashing" lectures at Princeton.

11.. It is interesting to note that Sosa makes such a move in "The Place of Truth in Epistemology," in Zagzebski and DePaul, eds., *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (New York: Oxford University Press, forthcoming 2003).

12.. Keith DeRose, "Solving the Skeptical Problem," *The Philosophical Review* 104 (1995), pp. 1-52.

13.. For a comprehensive account of the issues involved in this problem, see John Hawthorne's forthcoming book *Knowledge and Lotteries*, (Oxford: Oxford University Press).

14. The case here first appeared in Keith Lehrer and Thomas D. Paxson, Jr., "Knowledge: Undefeated Justified True Belief," *Journal of Philosophy* 66 (1969), pp. 225-237.