

Nozickian Epistemology and the Value of Knowledge

Epistemology involves reflection on cognitive achievements of a theoretical sort, the primary example of which is knowledge. The history of epistemology has focused primarily on questions concerning the nature and extent of human knowledge, but there is another question in epistemology's history that is equally important though little addressed, which is the question of the value of knowledge. Knowledge, we assume is valuable, and more valuable than mere belief, mere true belief, and even more valuable than mere justified true belief. The problem I will here address is how to explain this value.

The question of the value of knowledge arises first in Plato's *Meno*, where Socrates queries Meno concerning the value of knowledge over that of true belief. Meno first proposes that knowledge is more valuable because of the practical benefits to which it contributed, but upon seeing counterexamples to this claim, Meno entertains two thoughts. First, he wonders whether knowledge really is valuable, and second, he ponders whether knowledge is different from true belief. Socrates does not pursue the first issue but devotes some attention to the second, citing as metaphor the statues of Daedalus. The moral of the story, according to Socrates, is something he claims to know: that knowledge is different from true opinion.

The direction of the discussion leaves open the possibility that Socrates is willing to abandon the idea that knowledge is more valuable than true belief, but such a path should not be viewed as initially attractive (and it is surely implausible to think that Socrates or Plato would ever propose such an idea). Knowledge may not be all that we legitimately seek from an intellectual point of view, but it ought to take powerful arguments to move us to give up on the

idea that knowledge is one such thing. Another way to put this point is that there is a presumption in favor of the idea that knowledge is valuable, and any adequate epistemology should, in the absence of countervailing considerations, be able to provide a good answer to Socrates's questions. The ideal would be for one's epistemology to contain an explanation or account of the value of knowledge, but at the very least, a good epistemology should not (without argument) rule out the value of knowledge or imply that it is all-things-considered very unlikely to be valuable.

The difficulty is that it is not an easy task to answer the Socratic queries concerning the value of knowledge. In section 1 below, I explain the difficulty. Once we see what is needed for an adequate explanation of the value of knowledge, I turn in section 2 to the merits of Nozickian epistemology for addressing the problem of the value of knowledge. In section 3, however, I argue that such hope disappoints. Truth-tracking epistemology provides is not able to navigate successfully between the dual difficulties of explaining the nature and extent of knowledge on the one hand and the value of knowledge on the other.

I. The Swamping Problem

Let us begin by asking with Socrates what is the difference in value between knowledge and true belief. Though the possible answers here are manifold, two come to mind immediately. The first is Meno's pragmatic theory: knowledge is valuable as a means to getting what you want, whether it is getting to Larissa or becoming a millionaire. Socrates' reply undermines this account, for true belief works every bit as well for getting what you want as does knowledge—if

you want to get to Larissa, a guide with true beliefs will work just as well as a guide with knowledge. The second, and more plausible, attempt proceeds in terms of the constituents of knowledge: knowledge is more valuable than true belief because it involves additional elements which are themselves valuable.

Thus enter into the picture concepts such as justification, certainty, reliability, infallibility, incorrigibility, having the right to be sure, being probably true, and other such concepts. What is distinctive about each of these properties, it is claimed, is that they are valuable properties for a belief to have: in each case, the value of a belief with that property exceeds the value of a belief lacking that property.

These properties can be divided into two groups: those that are conceptually stronger than truth, and those that are conceptually weaker than truth. Among the former category are concepts that imply, but are not implied by true belief: infallibility, incorrigibility, and (objective) certainty. Among the latter category are concepts compatible with false belief: reliability, justification, having the right to be sure, probability of truth, and subjective certainty.

Concepts in the latter category face severe difficulty in addressing Socrates's concerns about the value of knowledge. Since each such concept is weaker than truth, each must find an answer to what I will call the swamping problem, which can be illustrated most clearly using the concept of probability. The difficulty encountered by appeals to this concept arises from the fact that the value of truth and the value of probability interact in a certain way. Socrates's question concerns the value of knowledge over that of true belief, so the proposed theory must hold that true beliefs that are likely to be true are more valuable than true beliefs that are not likely to be true. This value judgement has little to commend it, however. Once we have stipulated that a

belief is true, adding that it is also likely to be true fails to enhance its value. Beautiful people are not more beautiful because they have beautiful parents; a good-tasting apple tastes no better because it comes from a tree that usually produces good-tasting apples; and well-functioning furniture functions no better because it is made by a factory likely to produce such. In general, where P is a good-making characteristic of a thing T, T does not become more valuable by being likely to have P when T already has P.

We should be careful not to confuse this point with a related, but distinct, issue. The claim here is not the claim that likelihood of truth is not valuable, nor that a belief likely to be true is not more valuable than one that is unlikely to be true. These claims, I assume, are true, but they are not the point in question here. Instead, the view I am arguing is false is that a *true* belief is not more valuable in virtue of being likely to be true.

The above examples may be too cryptic to convince, so consider the following more careful example. Suppose I like chocolate and wish to have some within the next few minutes. Being new to the area, I consult the web, searching for places that sell chocolate and are within walking distance of my present location. My search turns up three sites. The first site is a list of places that sell chocolate and are within walking distance. The second site is a list of places that are likely to sell chocolate and are within walking distance. The third list is the intersection of the first two: it is a list of places that both sell chocolate and are likely to sell chocolate and are within walking distance. Given that the only relevant interest at hand is my interest in having chocolate soon, two points are true. First, the second list is inferior to the other two lists: given my interest in chocolate, a list of places that actually sell it is preferable to a list of places that are likely to sell it. Second, and more important, the first list is every bit as good as the third:

however likelihood of selling chocolate is measured, it adds nothing of interest in the present context to know of a place that actually sells chocolate that it is also likely to sell chocolate.

This phenomenon is the swamping problem. The problem is that when a property is valuable by being instrumentally related to another property, the value of the instrumental property is swamped by the presence of that for which it is instrumental. This fact alone does not undermine every attempt to explain the value of knowledge over true belief in terms of some property conceptually weaker than truth, for a property conceptually weaker than truth might be valuable on grounds unrelated to its relationship to truth. This possibility, however, offers little hope, for it immediately faces the problem of the truth connection. That problem is this: any property that distinguishes knowledge from true belief must bear some significant connection to the truth of belief, on pain of being irrelevant to the distinction between knowledge and true belief. A common expression of the truth connection is that whatever closes the gap between true belief and knowledge must make for likelihood of truth. If so, however, any property conceptually weaker than truth which has any hope of satisfying the truth connection will have to be conceived as instrumentally related to truth. As soon as this feature is included, the swamping problem appears.

Let me make clear here that I am not endorsing this entire line of reasoning. Instead, I only wish to display its plausibility. In my view, such reasoning contains important mistakes, but this is not the place for such a discussion. Instead, I only want to point out the plausibility of this line of reasoning in order to explain why it will be difficult to sustain the theory that properties weaker than truth and instrumentally related to it are up to the task of answering Socrates's queries.

As soon as this point is appreciated, it is easy to see why a vast majority of the history of epistemology focuses on the other category, the category including concepts stronger than truth. Not that the problem of the value of knowledge is the only, or even the primary, force behind this predilection, for such a thesis ignores the pervasive and dominant influence of skepticism on the history of epistemology. Yet, a point remains, even if we grant the primary influence of skepticism, for no property conceptually weaker than truth and only instrumentally related to it appears to have much hope of escaping the swamping problem.

This latter point is especially important once we move into the post-Moorean period in which the falsity of skepticism is assumed. Even if we make this assumption, the problem of the value of knowledge still motivates us to look beyond the category of concepts conceptually weaker than truth in attempting to distinguish knowledge from true belief. It is here that Nozickian epistemology is instructive and helpful.

II. Nozickian Epistemology and the Value of Knowledge

If we adopt the Moorean assumption that skepticism is false, but also heed the lesson above concerning the swamping problem, we will be led to favor accounts of knowledge that clarify the difference between knowledge and true belief in terms of properties conceptually stronger than truth. Such an approach, however, will make it difficult to keep skepticism at bay for long. Consider the property of infallibility, for example. Since it is conceptually stronger than truth, an infallible true belief is clearly more valuable than a mere true belief—it is better to be incapable of error regarding a particular subject than merely to avoid error. The problem is

that if knowledge requires infallibility, then we know (nearly) nothing. Skepticism reenters the picture, not because the skeptic gets to dictate the agenda of epistemology, but because the problem of the value of knowledge leads us to focus on concepts stronger than truth, and infallibility is a paradigm example of such a property.

The same result occurs when we consider other properties that are assumed to be conceptually stronger than truth, such as Descartes' concept of metaphysical certainty. False, metaphysically certain beliefs are not possible, according to Descartes, and because they are not, metaphysical certainty is conceptually stronger than truth. But it is so strong that it is doubtful that anything but the most modest of successes can be achieved in avoiding skepticism on a theory of knowledge maintaining that such certainty is required for knowledge.

It is here that Nozickian epistemology has an advantage. According to Nozick, knowledge is a factive doxastic state that is counterfactual supporting, so that if one knows that it is raining, then if it weren't raining one would not believe that it is raining and if it were raining one would believe that it is raining. Nozick's truth-tracking theory thus presents a condition for knowledge that is stronger than truth in the following sense: a condition C is stronger than truth if and only if (a) it is possible to have a true belief even though one's belief does not satisfy condition C, but (b) necessarily, when a belief satisfies condition C, it is a true belief. To see that Nozick's truth-tracking condition is stronger than truth, consider a particular example. Suppose that one believes that it is raining. Clause (a) is obviously true, since one could be doxastically promiscuous, being convinced by any sincere utterance by a completely untrustworthy source. To see that clause (b) is true, we assume that the belief is false. If the belief is false, the first counterfactual, usually called the sensitivity claim, will be false as well: it would be false that if it

weren't raining, one wouldn't believe that it is raining (since by hypothesis it isn't raining and one believes that it is raining).

Because truth-tracking is stronger than truth it offers hope of explaining the value of knowledge and not only its nature. It implies that knowledge is more valuable than mere true belief since being a truth-tracker is better than merely being a true believer. Thus even in the presence of the truth of a belief, adding the property of truth-tracking adds something to the composite that is not swamped by the value of truth itself.

III. Sensitivity and Safety

Nozickian epistemology thus passes a test for an adequate theory of knowledge not passed by many popular contemporary approaches to knowledge, such as process reliabilism of the sort Goldman proposes¹ (which, when considered from the perspective of the question of the value of knowledge, offers nothing more than the probabilistic theory used to explain the swamping problem). The harder question is whether this theoretical virtue can co-exist with a defense of the adequacy of such a Nozickian approach to the nature of knowledge. I will argue that the simpler versions of Nozickian epistemology are not promising at all, and that refinements that offer some hope of a more adequate account of the nature of knowledge force the loss of this theoretical virtue regarding the value of knowledge.

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 transparency 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 transparency condition have been popular in recent epistemology, and for good reason. The transparency condition would limit what we know to things that wouldn't be true unawares. Yet, there is good reason to think that much of what we know requires attention on our part. 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.

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.

Equally instructive is the fact that neither of these conditions yield a condition on knowledge stronger than truth. A false belief can satisfy both either the idealism condition or the transparency condition (or both), implying that a belief can meet these conditions without being a true belief. Given our starting point that identifies the promise of Nozickian epistemology in terms of an ability to explain the value of knowledge in virtue of adverting to conditions for knowledge that are stronger than truth, this failure gives a reason independent of those above for

rejecting counterfactual approaches that focus on these two conditions.

Such considerations explain the focus of recent epistemology on the other two conditionals, the safety and sensitivity conditions. Nozick uses the sensitivity condition, while epistemologists such as Sosa and Williamson prefer the safety condition. Both approaches are Nozickian in the sense that they distinguish knowledge from true belief on the basis of subjunctive conditionals. Moreover, the use of such conditionals offers some promise of accounting for the value of knowledge over that of true belief, but this promise is hollow if these approaches provide inadequate accounts of the nature of knowledge. One ground on which such approaches have been criticized concerns the question of whether knowledge is closed under known deduction. Exploring this issue will reveal that both conditions encounter problems on this score.

A. Closure

A theory of knowledge respects closure when it adopts the some version of the idea that we know the logical consequences of what we know. This rough closure claim is certainly false, otherwise we'd all be logically and mathematically omniscient. Theories that wish to respect closure defend the view by restricting this rough form of the view. One way to restrict it is to hold that when we know p and know p entails q , then we know (or are in a position to know) q .

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 in order to explain why a claim to knowledge is mistaken. When identifying a

particular SUV as a Ford Explorer, my assertion may be 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 makes 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 involves some sort of closure principle about knowledge, one that I accepted in the context when I reply that I saw the Ford logo on the front hood. Closure is initially attractive precisely because it provides a plausible explanation of the type of discussion this example illustrates.

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. For example, on Nozick’s view, it is possible to know that some object is a hand, know that if that object is a hand then no deception by a Cartesian evil demon is occurring, and yet not know that no such deception is occurring. Not many have been impressed with this aspect of Nozick’s view, and some have claimed that replacing sensitivity conditions with safety conditions gives hope for upholding closure.²

Consider the following two arguments:

1. Bp	1. Bp
2. p	2. p
3. p entails q	3. p entails q
4. $\sim p \square \rightarrow \sim Bp$	4. $Bp \square \rightarrow p$
5. So, $\sim q \square \rightarrow \sim Bq$	5. So, $Bq \square \rightarrow q$

The first two premises of each argument are not needed for the points that follow, but since the

context is one of attempting to ascertain the difference between knowledge and true belief, I include them to serve as a reminder of the context. We know that the first argument fails, on the basis of the example given above about hands and Cartesian demons. If I had no hands, I wouldn't believe that I do, so the belief that I have hands is sensitive to the facts. 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.

Notice that this example provides no counterexample to safety, thereby lending support to the idea that safety might preserve closure even though sensitivity does not. If one were to believe one has hands, one would have hands, so this belief is safe. Furthermore, if one were to believe that there is no demon, there would be no demon, 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.

Still, concluding that safety is closed under entailment would be incautious. First, note that it would be 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 : $p, Bp, q, \sim Bq, p$ entails q .

This stipulation guarantees the truth of the first three premises of the above arguments. Further stipulate that in all worlds \mathbf{W} most similar to α , we have:

True-in-W: $\sim p, \sim Bp, \sim q, Bq$.

This stipulation does not conflict with any reasonable interpretation of the closeness relation on worlds, so it is a model for the above two arguments. On this model, the premises of both arguments are all 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 transposition fails for counterfactual conditionals, but that fact is compatible with such transpositives nearly always having the same truth value. So it shouldn't surprise us if the transpositives 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 transposition 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. The residents of this locale have changed the ordinary landscape in the following way. They have made a decision for each barn whether to replace it with a fake barn or paint it green. We may imagine the decision-making to proceed in any number of ways, but let us suppose that they draw a card from a fair deck, and they replace if the card is black and paint if the card is red. Suppose that, in this area, 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, for it is a matter of chance whether he’d be correct (depending on whether the drawn card is black or red).³ 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. 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	3. p entails q
2. $Bp \square \rightarrow p$	4. $Bp \square \rightarrow p$
5. So, $Bp \square \rightarrow q$	5. So, $Bq \square \rightarrow q$

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 a 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, why should the theory rule out this kind of reliable indicator as suitable in the account of knowledge?

Reliable indicator theorists may respond by claiming that the reliable indicator view is

simply the deeper explanation of the truth of the safety view, in the following way. They may say, “If something is a reliable indicator of p, then believing p is also a reliable indicator of p, so reliable indication explains what is correct about the safety view.” The idea behind such a claim is to keep the counterfactual relationship between beliefs and propositions central to the view, so as to maintain in some form an analogy between the safety view and the reliable indication view. The problem is that 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
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.”)

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.

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.

From the perspective of the problem of the value of knowledge, this latter point is easy to appreciate. The characteristic feature of a Nozickian approach to the value of knowledge centers

on finding a condition for knowledge that is stronger than truth. It is in virtue of this feature that transparency and idealism conditions are not promising counterfactual approaches to the question of the value of knowledge. It is noteworthy, then, that a counterfactual reliable indication theory fails to offer a condition for knowledge that is strictly stronger than truth. To offer such a condition requires offering a condition for knowing p which implies that a belief with content p satisfying this condition must be true. When the reliable indicator is the belief itself, as is the case with the safety condition, this condition is clearly met: the result is obtained by an application of counterfactual *modus ponens*. No such inference is available, however, for the more general reliable indication theory. The only way such an inference could work is if the identity conditions on belief included its grounding in terms of the reliable indicator in question, but that view of belief is implausible. So the conclusion to draw is that the reliable indication theory is as impotent in addressing the problem of the value of knowledge as are the transparency and idealism counterfactual conditions dismissed earlier.

B. Probability and Closure

A natural move at this point is to wonder whether closure 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 the claim in question is usually true in the presence of belief,⁴ and one might say that a belief is sensitive when one would typically not hold the belief if the claim were false.

Both of these claims contain a scope ambiguity, 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 fares 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. Bp	1. Bp
2. p	2. p
3. p entails q	3. p entails q
4. $\Pr(\sim p \Box \rightarrow \sim Bp)$	4. $\Pr(Bp \Box \rightarrow p)$
5. So, $\Pr(\sim q \Box \rightarrow \sim Bq)$	5. So, $\Pr(Bq \Box \rightarrow q)$

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 and nomological ones, namely that the truth of the conclusion in face of the truth of the premises will be contingent at best.

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 logically necessary 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 a fair 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.⁵ In the case of a relatively large lottery where one believes that one's ticket will lose, one's belief is, if not safe, at least probably safe. The lesson is that probabilistic safety theorists will need some further condition to give an adequate account of the nature of knowledge.

Furthermore, the move to a probabilistic understanding of safety removes the primary virtue of counterfactual approaches to the nature of knowledge. Recall the initial attraction of such theories: the capacity to explain the value of knowledge in virtue of citing a condition strictly stronger than truth. Once a probability operator is introduced, either of wide or narrow scope, the new condition is no longer strictly stronger than truth. In particular, no such condition guarantees that a belief that meets such a condition must be true. Believing a false claim is compatible both with it being likely that the belief counterfactually implies its content and with the belief counterfactually implying the likely truth of its content.

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. In the context of our discussion, the worst of these problems is that such a move eviscerates the safety approach of its promise of accounting for the value of knowledge.

IV. Deeper Problems for Safety and Sensitivity Conditions

Two responses might blunt the force of the above discussion of varieties of counterfactual approaches in epistemology. First, one might find some way of preserving closure for one of the

two views in some way I have not considered. Given the variety of considerations above, I think it is highly doubtful that any such approach can be found, but I've given no decisive proof against the idea. Second, even if no rescue is possible, it may be possible to preserve Nozick's central contention that denies closure while restricting the failures in such a way that the denial is more palatable. For example, it is hard to defend the view 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."⁶ 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. So perhaps there is a way to develop safety or sensitivity to reduce the kinds of failures of closure to those that are the least problematic kind, and to argue that the virtues of the approach are sufficient to outweigh the disvalue of this negative aspect of the theory.

Furthermore, I don't think the issue of closure gets to the heart of the matter of what is troublesome about these approaches to the nature of knowledge, for recent epistemology has found it extremely difficult to formulate a correct account of closure.⁷ 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 are suitable only for explaining 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, as formulated in 2 and 3 above, 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.⁸ 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 will die.

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. 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 has 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 cannot rule 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 to account for the value of knowledge over that of true belief.

Conclusion

These results are deeply disappointing. We began with an intuitive account of how Nozickian approaches to epistemology might help us understand the value of knowledge over that of true belief. Nozickian approaches have a virtue not possessed by other contemporary theories of knowledge, for such accounts impose a condition on knowledge logically stronger than truth. This feature endows these theories with the same resources for explaining the value of knowledge over that of true belief as are possessed by historically prominent infallibilist theories of knowledge, while avoiding the nearly universal skepticism implied by the latter. Given this strong promise of Nozick's approach, it is regrettable that we must conclude it to be a failure. This failure raises a larger question that deserves careful investigation. That question is the question of whether any approach to the theory of knowledge can successfully explain the value of knowledge over that of its subparts (without encountering other problems such as implying the truth of wholesale skepticism). Our discussion can be read as providing significant

evidence for a negative answer to this question,⁹ thereby robbing the theory of knowledge of one argument for its centrality in a proper conception of the tasks of philosophy.¹⁰

Endnotes

1. See, e.g., Alvin Goldman, “What is Justified Belief?” in G. Pappas, ed., *Knowledge and Justification*, (Dordrecht: D. Reidel, 1979), pp. 1-25.
2. A clear and helpful discussion of these issues can be found in Steven Luper, “The Epistemic Closure Principle,” in the *Stanford Encyclopedia of Philosophy*.
3. 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.
4. 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).
5. There are contexts in which no criticism is warranted for an assertion that one knows that one ticket will not win (e.g., if one has to choose between waiting to receive the prize for winning and attending to some emergency). My view of such cases is that a careful account of the pragmatics/semantics distinction applied to the concept of knowledge will count such assertions as pragmatically acceptable even though semantically false.
6. Keith DeRose, “Solving the Skeptical Problem,” *The Philosophical Review* 104 (1995), pp. 1-52.
7. For a comprehensive account of the issues involved in this problem, see John Hawthorne’s forthcoming book *Knowledge and Lotteries*, (Oxford: Oxford University Press).
8. What I am claiming does not imply that there are no counterexamples to sensitivity in cases of perceptual knowledge, nor that there are no counterexamples to safety in statistical or inductive cases. I am only pointing out the natural fit between sensitivity considerations and perceptual knowledge, on the one hand, and safety considerations and statistical knowledge, on the other.
9. For further argument for this conclusion, see my *The Value of Knowledge and the Pursuit of Understanding*, (New York: Cambridge University Press, 2003).
10. For helpful discussion of these issues and comments on earlier drafts, I wish to thank Peter Markie, Matt McGrath, Peter Vallentyne, and Paul Weirich. I also wish to thank audiences at the University of Arkansas, St. Louis University, and North Carolina State University for useful and provocative discussion of these issues.