JONATHAN VOGEL

ARE THERE COUNTEREXAMPLES TO THE CLOSURE PRINCIPLE?

Very often, a person can't know a proposition without knowing various logical consequences of that proposition. So, for instance, if you know that your friend is wearing a yellow tie, you can't fail to know that your friend is wearing a tie, period. In this case, the relation of logical consequence is obvious. When the relation isn't obvious, a proposition you know may have a logical consequence you don't know — for example, a suitably obscure mathematical theorem. In light of these considerations, it seems plausible to hold that if a person knows a given proposition, that person must also know any logical consequence of that proposition which he or she recognizes as such. Putting it differently, we might say that knowledge is closed under known logical implication.[1]

The problem of skepticism about the external world gives this epistemic principle (hereafter, the "Closure Principle") a special interest. When the skeptic argues that we have no knowledge of the world because we don't know that we aren't massively deceived in some way, he or she appears to assume that knowledge has the closure property. But if it is possible to find clear examples demonstrating that closure sometimes fails, a crucial piece of support for skepticism will be removed. The purpose of this paper is to show that even the strongest apparent counterexamples to closure don't hold up under scrutiny. To that extent, the problem of skepticism is still with us.

I DRETSKE'S ZEBRA CASE

In a widely read paper, Fred Dretske offered an intriguing example which is meant to show that the Closure Principle is invalid. It is worthwhile to quote Dretske's discussion at length:

> You take your son to the zoo, see several zebras, and when questioned by your son, tell him they are zebras. Do you know they are zebras? Well, most of us would have little hesitation in saying that we did know this. We know what zebras look like, and, besides, this is the city zoo and the animals are in a pen clearly marked "Zebras." Yet, something's being a zebra implies that it is not a mule and, in particular, not a mule cleverly disguised by the zoo authorities to look like a zebra. Do you know that these animals are not mules cleverly disguised by the zoo authorities to look like zebras? If you are tempted to say "Yes" to this question, think a moment about what reasons you have, what evidence you can produce in favor of this claim. The evidence you *had* for thinking them zebras has been effectively neutralized, since it does not count toward their *not* being mules cleverly disguised to look like zebras. You have some general uniformities on which you rely, regularities to which you give expression by such remarks as "That isn't very likely" or "Why should the zoo authorities do that?" Granted, the hypothesis (if we may call it that) is not very plausible, given what we know about people and zoos. But the question here is not whether this alternative is plausible, not whether it is more or less plausible than that there are real zebras in the pen, but whether *you know* that this alternative hypothesis is false. I don't think you do[2].

According to Dretske, the Zebra Case is a counterexample to closure because you know (a) the animals in the pen are zebras, but don't know a clear logical consequence of (a), namely,(b) the animals in the pen aren't cleverly disguised mules. I find this description of the situation implausible. Given what Dretske has said in laying out the example, I think it is more reasonable to conclude that if you know (a) you know (b) as well, and closure is preserved after all.

The reason you know that an animal in the pen is not a disguised mule (if you do know it's a zebra) is that you have a true belief to that effect backed up by good evidence. That evidence includes background information about the nature and function of zoos. You know that zoos generally exhibit genuine specimens, and that it would be a great deal of trouble to disguise a mule and to substitute it for a zebra. Only under the most unlikely and bizarre circumstances, if at all, would such a substitution be made, and there is no reason whatsoever to think that any such circumstances obtain. If you did feel there was a chance that a switch had been made, you would have reason to doubt that the animal you see is a *zebra*. You would not, then, know that it is a zebra, contrary to what was assumed.

Dretske's motivations for denying that you know you aren't seeing a disguised mule are not fully clear. He himself grants that the "hypothesis" that the animal is really a mule is "not very plausible", yet adds

> But the question here is not whether this alternative is plausible, not whether it is more or less plausible than that there are real zebras in the pen, but whether *you know* that this alternative is false.[3]

One might have thought that if a belief is much more plausible than its denial, a person would be justified in accepting that belief. And, then, barring Gettier–like complications, that person's belief, if true, would be knowledge.[4]

Perhaps Dretske's point is this: When you look at the pen where the animal is, you have evidence that there is a zebra there, namely that the animal looks like a zebra. Your visual evidence does not, though, give any support to your belief that the animal you are seeing *isn't* a disguised mule. For, if it were a disguised mule, your visual experience would be just as it is. As Dretske says, "The evidence you *had* for thinking them zebras has been effectively neutralized, since it does not count toward their *not* being mules cleverly disguised to look like zebras".[5] The upshot is that you do know there is a zebra, since you have a true belief to that effect supported by evidence. You do *not* know that the animal isn't a disguised mule, since your belief in this case is true but not supported by available evidence. So, you know the first proposition, but don't know its clear logical consequence.

I indicated above why I think this analysis is incorrect. Your background knowledge does give you justification for denying that the animal is a mule, so you know that it isn't one. Still, it may appear that the possibility of failure for the Closure Principle arises out of the situation as I described it. It seems that the usual adequate evidence for the claim "It's a zebra" (i.e. visual evidence) is different from the background evidence which supports "It's not a cleverly disguised mule." If so, you could conceivably be in a position where you had the visual evidence and knew there was a zebra, but lacked the background knowledge, and hence didn't know there wasn't a disguised mule. In such circumstances, the Closure Principle would face a counterexample.

To my mind, this appraisal is based on an overly atomistic conception of evidence and justification. Your belief that the animal at the zoo is a zebra is justified in part by your visual evidence, but it is *also* supported by the background information that counts against the animal's being a disguised mule. By itself, the visual evidence wouldn't be sufficient to give you knowledge that there is a zebra. To see this, consider a case where the proper background knowledge is lacking. Imagine that you are driving through ranchland out West and for some reason or other stop by the roadside. Across the way you see a black and white striped equine creature tranquilly grazing in its pen. In a situation of this sort, it seems to me, it is far from clear that you could know the animal before you to be a zebra, even though it looks just as much like a zebra as the animal in the zoo does. The difference here is that you have no applicable background information which makes it more likely that a zebra–like animal really is a zebra rather than an oddly colored mule. So, even back at the zoo, your justification that what you see is a zebra depends on background information — just as the justification for your denial that it's a disguised mule would so depend.[6] There is no discrepancy here which provides grounds for thinking that the Closure Principle is false.

One might object that the defense of closure just given makes unrealistically high demands so far as evidence is concerned. A young child at the zoo, seeing an animal that resembles an illustration in a picture book might point and happily say "Zebra!". Despite the fact that the child knows nothing about how zoos work, doesn't that child know the animal is a zebra? The issues here are complex, but there are various reasons not to take this objection as decisive. First, even if it is granted that the child knows in the full sense that the animal is a zebra, if he or she isn't capable of drawing the inference about disguised mules, the child's case doesn't bear on the validity of the Closure Principle. Moreover, it's unclear that, under the circumstances, the child really ought to be described as knowing that the animal is a *zebra*. Suppose that the child can't conceptually distinguish between 'looks like an zebra' and 'is a zebra'. Perhaps the child knows only that the animal it sees looks like a zebra, and wouldn't know that the animal is a zebra without acquiring further conceptual resources and information.[7]

II CAR THEFT CASES

I have maintained that Dretske's Zebra Case does not furnish a counterexample to the Closure Principle. But what I have said so far bears largely on the particular details of the case as Dretske sets it up. His remarks point towards the formulation of examples which cannot be treated so straightforwardly. I call these "Car Theft Cases", for reasons which will become clear in a moment. It may be, in fact, that the Zebra Case properly understood is one of these.

Suppose you own a car which you parked a few hours ago on a side street in a major metropolitan area. You remember clearly where you left it. Do you know where your car is? We are inclined to say that you do. Now it is true that every day hundreds of cars are stolen in the major cities of the United States. Do you know that your car has not been stolen? Many people have the intuition that you would not know that. If this intuition is combined with the previous one, then it seems that the closure principle is

violated. That is: You know the proposition (p) 'My car is now parked on (say) Avenue A'. You also know that that proposition entails (q) 'My car has not been stolen and driven away from where it was parked'. Yet, it seems, you do not know q, despite the fact that it is for you a clear logical consequence of p, which you do know. Since, in this instance, you (apparently) fail to know a clear logical consequence of a proposition you do know, the Closure Principle is (apparently) violated.

This example turns on a rather unusual feature of the clear logical consequence q. Given your evidence, that proposition is much more probable than not, and it is at least as likely to be true as p is. To that extent, it seems as though you should be as justified in believing q as you are in believing p. Nevertheless, even though your belief that p, if true, may be knowledge, your belief that q, if true, is not. You do not know that your car hasn't been stolen by someone and driven away, despite the high probability that your belief to that effect is true.

In this respect, your belief that q resembles someone's belief that a ticket, which he holds, will not win a fair lottery. No matter how high the odds that the ticket will not win, it strikes us that the ticket-holder doesn't *know* that his ticket will not win. In fact, the analogy between a subject's belief about holding a losing lottery ticket and one's belief that one's car has not been stolen goes even further than this and is quite illuminating.

A number of features of a lottery situation are especially relevant here. First, although winning a lottery on a particular ticket is unlikely or improbable, it would not be *abnormal* in some intuitive sense, for it to turn out that the ticket one holds happens to be a winner. Second, even though the weight of the evidence is certainly against any particular ticket's winning, there is still some *statistical* evidence in favor of the proposition that a certain particular ticket will win, i. e. there is some (small) reason to think a particular ticket-holder will win.[8]

A third important consideration is that, with respect to its chances of winning the lottery, each ticket is indistinguishable from every other one. So, any reason you have for thinking that your particular ticket will lose would be an equally good reason for believing of any other ticket in the lottery that it, too, will lose. Under these circumstances, it would be arbitrary to believe of some tickets (including your own) but not others that they will not win. So, if you are consistent rather than arbitrary, and you do conclude on the basis of the evidence available that *your* ticket will not win, you will conclude the same of every other lottery ticket. Nevertheless, you hold the belief that some ticket or other will win. On pain of arbitrariness, then, it seems that you can't justifiably hold both that your ticket will lose *and* that some ticket will win. *A fortiori*, you can't know that your ticket will lose and that some ticket will win.[9]

Now, in certain important ways, one's epistemic situation with respect to the lottery is like one's epistemic situation in the Car Theft Case.[10] In effect, when you park your car in an area with an appreciable rate of auto theft, you enter a lottery in which cars are picked, essentially at random, to be stolen and driven away. Having your car stolen is the unfortunate counterpart to winning the lottery. And, just as one doesn't know that one will *not* have one's number come up in the lottery, it seems one doesn't know that one's number won't come up, so to speak, for car theft.

To be more particular, believing that your car won't be stolen is like believing you won't win the lottery, in the ways just canvassed. (1) If you park your car in an area with a high rate of car theft, an area where it is virtually certain that some car like yours will be stolen, it would not be abnormal for your car to be stolen. (2) In the Car Theft Case, your knowledge that there is a considerable amount of auto theft gives you some real statistical reason to think you car will be stolen.[11] (3) It would be arbitrary of you to believe that your car, but not all the others relevantly similar to it, won't be stolen. In general, if a person fails to know a proposition because of considerations like these, I will call the proposition not known a *lottery proposition*.

The point of this extended comparison of the lottery and the Car Theft Case has been to try to characterize a family of apparent counterexamples to the Closure Principle. The essential feature of these examples is that they are cases in which the clear logical consequence of a known proposition is itself a lottery proposition meeting the criteria just discussed. What makes the Zebra Case, in my opinion, a weaker potential counterexample to the Closure Principle than the Car Theft Case, is just the fact that the clear logical consequence of the Zebra Case is harder to see as a lottery proposition. First, it *would* be abnormal for a disguised mule to be in a zoo enclosure marked "Zebras". Second, as Dretske describes the example, it isn't apparent that you have any reason (statistical or otherwise) to think that there might be a disguised mule in the zebra pen. These two weaknesses are related to the third: it is difficult to see the presence of a disguised mule in the zebra pen as the outcome of any lottery-like process. That is, it is not as though you know that a disguised mule has been placed in some zebra pen in some zoo chosen at random. In *that* case, any reason you had for thinking that the animal you happen to see isn't the disguised mule would apply in every other situation. You would, then, have to conclude that no zoo had a disguised mule running around — in contradiction with what you know to be the case, viz. there is a disguised mule in some zoo somewhere. However, this kind of lottery element isn't present in the Zebra Case as Dretske described it. So, it is unclear why, as Dretske maintains, you do not know that the striped animal before you isn't a disguised mule.[12]

III CAR THEFT CASES AND SKEPTICISM

I would like to turn now to the implications of the Car Theft Case. That case is supposed to count as a counterexample to the Closure Principle. For, in the Car Theft Case, you seem to know a proposition about where your car is, but you apparently fail to know another proposition which is a clear logical consequence of the first one. I will maintain below that taking the Car Theft Case in this fashion, as a counterexample to closure, is not the only, or the best way, to understand it. But, suppose that the Car Theft Case does stand as a counterexample to closure; does that really help us with the problem of skepticism?

The thought was that the Car Theft Case would show that closure isn't valid in general. Then the skeptic's reliance on that principle in the course of the argument from deception would be illegitimate, and the argument wouldn't go through. However, what the Car Theft Case really shows about the Closure Principle, if it shows anything at all, is that that principle is invalid when the clear logical consequence involved is a lottery proposition *with the features mentioned above*. The Car Theft Case gives us no reason to think that closure fails to hold for clear logical consequences which don't satisfy those criteria.

The question at this point is whether the clear logical consequence in the skeptic's argument is a lottery proposition in the specified sense. The clear logical consequence the skeptic invokes is something like 'I am not a brain in a vat thoroughly deceived by sinister neurophysiologists'. And this is clearly *not* a lottery proposition satisfying the three criteria having to do with abnormality, reliance on statistical evidence, and non–arbitrariness. Let me take these out of order. (1) If the skeptic's logical consequence were a lottery proposition, I would have to be an indistinguishable member of a class of subjects of which it is known that at least one member is a brain in a vat (making it arbitrary for me to believe that I'm not such a brain). This is hardly the case, since I don't know that there are any brains in vats anywhere. The lottery–like element which was crucial to the structure of the Car Theft Case is therefore lacking here. (2) Moreover, since there is no reason to think that some brains are put into vats as a matter of course, it might well be abnormal, in an intuitive sense, for someone to turn out to be a brain in a vat. (3) Finally, given (1), there is no basis for assigning a real, positive statistical probability to the proposition that someone is a brain in a vat.

The force of these observations is that the situation in which the skeptic invokes closure cannot easily be assimilated to situations like the Car Theft Case, in which there is some reason to think closure fails. Hence, the Car Theft Case as such gives little support to the claim that the Closure Principle fails when the skeptic appeals to it. This means that the Car Theft Case provides no convincing basis for rejecting the Deceiver Argument.

It may be that, if Cartesian skepticism is the issue, no more needs to be said about the Zebra Case or the Car Theft Case. I will, however, pursue the question of whether the Car Theft Case is a genuine counterexample to the Closure Principle. Aside from whatever intrinsic interest that question may have, it is worth seeing that the results strengthen, rather than weaken, the conclusion that these examples do not undercut skepticism.

IV THE INTERPRETATION OF INTUITIONS ABOUT THE PROBLEM CASES

The Car Theft Case and its analogues provide counterexamples to the Closure Principle if we take our intuitions about such cases at face–value. For, then, it seems that in the circumstances described, a person may know some proposition (e.g. 'My car is on Avenue A, where I parked it') yet not know a clear logical consequence of that proposition (e.g. 'My car hasn't been stolen and driven away from where it was parked'). It's worth noting, though, that some additional reactions people have suggest that closure is preserved in these situations after all. Often, when faced with the possibility that their cars might have been stolen, people withdraw, at least temporarily, their initial claims to know where their cars are. Such a response is just what the Closure Principle would require.

Now, I think it must be admitted that the intuitions we have here are weak. It would be difficult to find decisive support for closure in the tendency people have to change their minds in the way just mentioned. Still, the fact that the Closure Principle seems to be respected to the extent that it is provides a motivation for analyzing that case in a way that doesn't presuppose the failure of closure.

The problem facing any such analysis is to accommodate or discredit the intuitions that produce the impression of closure failure in the first place. Those are the intuitions which lead us to say, first, that a person, under certain circumstances, would know some proposition, and, second, that the person doesn't know a clear logical consequence of that proposition. One way of trying to reconcile these intuitions with closure is to argue that some kind of shift takes place between these responses. The claim would then be that, for no *fixed* set of circumstances, do we regard a subject as knowing a proposition while failing to know one of its clear logical consequences.

Certain psychological studies provide independent reasons to believe that a shift of this kind takes place. These studies concern people's attitudes towards improbable events. They are relevant to the Car Theft Case because of the essential role played in that case by the unlikely possibility that your car has been stolen. If closure does fail here, it is because the possibility of theft, though highly improbable, undercuts the claim that you *know* that your car hasn't been stolen, even while that possibility somehow leaves intact your knowing that your car is at a certain spot. In the studies mentioned, it has been found that people may treat improbable events either as likelier than they really are or as having essentially no chance of occurring. Moreover, these assessments are unstable, and subjects can easily be influenced to grant a possibility more weight than otherwise, if that possibility is made salient to them.[13]

Such psychological considerations provide an explanation for our intuitions about the Car Theft Case. Initially and generally, in evaluating the knowledge claims in that case, we treat the chance of your car's being stolen as essentially zero. You can, then, be as sure as you need to be that your car is where you left it; you are fully justified in that belief. Thus, we are likely to say without hesitation that in the situation described you know where your car is. Later, however, when we dwell on the rate of car theft, the chance of your car's having been stolen is lent more weight. Given a (now) significant possibility that you may be wrong in believing that your car hasn't been stolen, we are no longer prepared to say that you *know* it hasn't been stolen. And, viewing the situation in this light, giving weight to the chance that the car isn't where you left it, we may be inclined to go on to say that you don't know where the car is after all. That is, there seems to be a motivation to deny your initial knowledge claim in a set of circumstances where you cannot claim to know a clear logical consequence of what you thought you knew. In that way, the Closure Principle is respected.

In short, the fact that at one time we would say that you know the location of your car, and that *shortly thereafter* we might say that you don't know your car hasn't been stolen, does not establish the invalidity of the Closure Principle. For, it may be that at no *one* time do we affirm that you know something yet fail to know one of its clear logical consequences. It is doubtful, then, that the Car Theft Case, when properly understood, provides a counterexample to the Closure Principle.

I have suggested that the anomalous character of our intuitions about the Car Theft Case may be due to some kind of epistemically important shift rather than to closure failure. My conjecture has been that the shift is a change in a probability assignment, but other mechanisms may be at work instead. An alternative explanation of our intuitions is that we are somehow induced to shift our sense of the degree of assurance knowledge requires. Thus, our estimation of the chance the subject could be wrong because of car theft would remain constant, but we would change our minds as to whether

knowledge is consistent with that level of epistemic risk. There are still other forms the shift could take. It might even be that the movement in the Car Theft–type situations is between wholly distinct notions of knowledge embodying different sets of necessary and sufficient conditions.

For my purposes, the details of what actually occurs are relatively unimportant. The main point I wish to make is that there are explanations other than closure failure for our intuitions about the Car Theft Cases.[14] Or, to put it differently, a straightforward appeal to those intuitions is insufficient to establish that the Closure Principle does not hold without restriction.

V THE PROBLEM OF SEMI-SKEPTICISM

I have just argued that a simple inspection of our intuitions about the Car Theft Case does not conclusively refute the Closure Principle. The advocate of closure can claim that the Closure Principle only *appears* to fail, as the result of an epistemically important switch that takes place in the course of our thinking about the example. However, a claim of this sort leaves open what a subject, in fact, does and doesn't know in Car Theft–type situations. The Closure Principle faces a strong objection to the effect that it is incompatible with any acceptable account of what is known in Car Theft Cases.

If closure holds, and some uniform standard of knowledge applies across the board, either you don't know where your car is, or you do know that it hasn't been stolen. The latter claim seems hard to sustain. This impression is strengthened by the similarity between the Car Theft Case and a real lottery situation. Knowing that your car hasn't been stolen would be, in the ways I've mentioned, like knowing someone will lose a fair lottery. And *that* seems like the sort of thing one doesn't know. So, given the untenability of saying that you know your car hasn't been stolen, the Closure Principle will require that, contrary to what we might have thought, you don't know where your car is.

This result seems unwelcome, and things worsen quickly. It turns out that, of the propositions about the external world which we take ourselves to know, a great many entail lottery propositions as in the Car Theft Case. (The propositions with these consequences are, specifically, propositions about the current state of the world beyond our immediate environments). To see the range of Car Theft–type cases consider some other examples:

Bush Case:

> Q. Do you know who the current President of the United States is?
> A. Yes, it's George Bush.
> Q. Do you know that Bush hasn't had a fatal heart attack in the last five minutes?
> A. No.

Luncheonette Case:

> Q. Do you know where I can get a good hamburger?
> A. Yes, there's a luncheonette several blocks from here.
> Q. Do you know that a fire hasn't just broken out there?
> A. No.

Meteorite Case:

> Q. Do you know what stands at the mouth of San Francisco Bay?
> A. Yes, the Bay is spanned by the Golden Gate Bridge.
> Q. Do you know that the Bridge wasn't just demolished by a falling meteorite?
> A. No.

It's apparent that variations on these cases can be constructed for any number of propositions about people, things, or activities. That is to say, all the propositions about such matters, which we take ourselves to know, entail lottery propositions which, it seems, we do not know. If closure holds, along with the intuition that we do not in fact know the clear logical consequences in question, the result is that we have a great deal less knowledge of the world than we had supposed. In other words, the Closure Principle leads, even without the argument from deception, to a fairly strong and unpalatable semi–skepticism. The case against closure appears that much the stronger.

But does the threat of semi–skepticism really count against the Closure Principle? The key idea here is that there is supposed to be some feature which the lottery propositions in Car Theft Cases share with propositions about genuine lotteries, in virtue of which we can't be correctly described as knowing those propositions. What is that feature? One answer is that, because of the statistical probability that your ticket may win in a genuine lottery, there is a "real" possibility of error in believing that you will lose. In other words, the crucial belief in these circumstances lacks a kind of certainty, and hence can't count as knowledge.[15] Similarly, the lottery propositions which figure in Car Theft Cases are such that a "real" possibility exists that they are false. Since, therefore, the subject can't be certain of the truth of these lottery propositions, the subject can't know them. By the Closure Principle, it would follow that the subject can't have knowledge of the propositions which he knows to entail those lottery propositions. This would result, as we have seen, in a pervasive semi–skepticism.

The important thing to realize about this way of viewing matters is that it doesn't really justify concluding that the Closure Principle is invalid. For, according to the objection, the lesson of the genuine lottery examples is that a belief can't be knowledge if there is a "real", and not merely logical, possibility that the subject is wrong about it. If this is correct, then semi–skepticism follows *without* the Closure Principle. After all, there is a "real" possibility that, e. g. you may be wrong in believing that your car is at a certain spot; it is possible that your car has been stolen. The same point applies, *mutatis mutandis*, to any other Car Theft Case. So, perhaps, there is a legitimate epistemological problem in the threat of a semi–skepticism derived from a certainty requirement for knowledge. However, since rejecting closure won't avoid that problem, that problem doesn't provide a reason for denying the Closure Principle's validity.

On another way of analyzing the lottery examples, the unknowability in these contexts of propositions like 'My ticket will lose' is due to the arbitrariness of accepting

any proposition of that form. By analogy, in the Car Theft Case, you wouldn't know the proposition 'My car has not been stolen'; there is reason to think that some car or cars similar to yours will be stolen, and you have no non–arbitrary ground for believing that your car in particular won't be the one (or one of the ones) stolen. Once more, it looks as though *all* knowledge claims about lottery propositions in other Car Theft cases would be undercut by similar considerations. Then, semi–skepticism will be inevitable if closure holds.

Here again, though, I am inclined to think that there is no argument to be found against the Closure Principle as such. The analysis of the lottery effect now being entertained makes the following assumption: all other things being equal, it is unjustified to accept any member of a set of propositions *L*, such that the members of *L* are equiprobable and the subject knows (or has good reason to believe) that at least one member of *L* is false.[16]

It turns out that this principle is sufficient to establish semi–skepticism regardless of the validity of the Closure Principle. To see why this might be so, let's take the Car Theft Case as the basic model. The present attempt to attach the burden of semi–skepticism to the Closure Principle amounts to the claim that the non–arbitrariness requirement just stated defeats your claim to know the lottery proposition that your car hasn't been stolen — while it leaves intact your claim to know a proposition (i.e. 'My car is on Avenue A, where I parked it') clearly entailing that lottery proposition. But the entailing proposition is itself a member of a set of equiprobable propositions which, you have good reason to believe, contains at least one falsehood. That set contains, along with 'My car is on Avenue A, where I parked it', propositions like 'My neighbor's car is where he parked it', 'The postman's car is where he parked it', and so on. You may not be able to state all the members of the set explicitly, but you still have very good reason to think that there is such a set *L*. By the non–arbitrariness requirement, it would follow that you don't know the original proposition 'My car is on Avenue A, where I parked it'.[17]

The same line of thought would seem to apply to any case of the Car Theft–type where knowledge of a lottery proposition is blocked by the non–arbitrariness constraint. So, if the non–arbitrariness condition is strong enough to establish ignorance across the board for lottery propositions, it is also strong enough to establish ignorance of the propositions which, in Car Theft cases, entail the lottery propositions. That is to say, if the non–arbitrariness condition *plus* closure generates semi–skepticism, so too does the non–arbitrariness condition alone. Therefore, the opponent of closure cannot use that condition as the basis for an argument that the Closure Principle is invalid because *it* would lead to semi–skepticism.

The preceding discussion makes clearer what would be required in order to make the case against closure work. The critic of the Closure Principle has to identify some way in which beliefs in lottery propositions are epistemically defective, and this defect must not be shared by the mundane beliefs whose contents, in Car Theft cases, are known to entail those lottery propositions. It isn't easy to see what such a defect would be, if not the ones just considered.[18]

In this section, I have tried to show that our anomalous intuitions about Car Theft Cases and the related threat of semi–skepticism really have little to do with closure. No attempt has been made here to give a fully acceptable positive account of what really

known in these cases, and I suspect that such an account may not be available at all. For it may be that the Car Theft Cases together with the problem of semi–skepticism reflect deep–seated, unresolved conflicts in the way we think about knowledge.[19]

VI CAR THEFT CASES AND RELEVANT ALTERNATIVES

It is tempting to think that the omission of a positive account of what we know could be made good by adopting a version of the relevant alternatives approach to knowledge.[20] This approach promises all the advantages, without the defects, of the treatment just given. In my view, a turn to the relevant alternatives approach is not advisable, but the proposal is interesting and deserves consideration.

According to the relevant alternatives theorist, the demands for knowledge are restricted and contextual. On one version of the theory, *S* knows that *p* just in case *S* possesses evidence which counts against all relevant alternatives to *p*; on another formulation, *S* knows that *p* just in case *S* would be right about *p* over some class of relevant alternative situations. A major problem for the relevant alternatives approach is to explicate the crucial notion of relevance it invokes. Relevance of alternatives will vary according to the subject's situation; it may also (depending on the details of the theory) be determined by the content of the subject's belief and the context of attribution for the knowledge claim. If the standard of relevance obeys certain constraints, the relevant alternatives theory may be used to explain intuitions about the Car Theft Cases in a way that doesn't deny the validity of the Closure Principle.

How would this go? Suppose the facts are as described in the Car Theft Case. Initially, we operate with a standard of relevance according to which the possibility of Car Theft is too remote to be considered. At this point, the fact that you would be wrong about the location of your car, had it been stolen,[21] doesn't impair the claim that you know where your car is. Moreover, since the possibility of car theft is remote, that possibility doesn't undercut the claim that you know your car hasn't been stolen. Closure is maintained. What produces the impression to the contrary? When the possibility of car theft is explicitly raised, somehow a new, more generous standard of relevance is instated, according to which the possibility of car theft *is* relevant. By this standard, you know neither where your car is nor that it hasn't been stolen. Closure is still preserved, as before.[22]

There are several drawbacks to analyzing the Car Theft Cases in this fashion. First, the supposed virtue of the analysis is that it provides an account of what you would and wouldn't know in the circumstances given. But in giving such an account, the relevant alternatives theorist must say that, in some sense or from some standpoint, you would know that your car hasn't been stolen. This seems plainly wrong, and the intuition that it is wrong is just what makes it so hard to give an adequate treatment of the Car Theft Case and its analogues. The relevant alternatives approach really doesn't accommodate the body of our intuitions in an unforced, convincing way, contrary to what one might have hoped.

Let me turn to a further point. The relevant alternatives theorist hypothesizes that, in the problem cases, there is a shift in the standard of epistemic relevance. In the Car Theft Case specifically, the possibility of car theft is supposed to be, alternatively, too remote and not too remote to be relevant. It is natural to presume that "remoteness" here

is to be understood in probabilistic terms. Thus, at one time, the chance of car theft is treated as small enough to be ignored; later, in a more scrupulous frame of mind, we find even that little probability of error sufficient to undercut knowledge. Relevance, then, is a function of an alternative's probability.

This probabilistic criterion of relevance seems attractive, but it leads to trouble, especially if knowledge requires having evidence that excludes relevant alternatives. Suppose you know a proposition *k*. Let *l* be an alternative probable enough to be relevant to *k*, and let *m* be any other alternative to *k* which should count as *ir*relevant. Consider, in addition, the disjunction $(l \vee m)$, which is logically incompatible with *k*. This disjunction is at least as probable as its disjunct *l*, so it is probable enough to be relevant to your knowing *k*. Now, since $(l \vee m)$ is relevant to your knowing *k*, you have to have good evidence against it. That is to say, you have to have good evidence for the negation of $(l \vee m)$, namely the conjunction (*not–l* & *not–m*).

Why is this a problem? If you have good evidence for (*not–l* & *not–m*), you presumably have good evidence for *not–m alone*.[23] Thus, your being in this favorable position with respect to *not–m* is a condition for your knowing *k*. So, *m isn't* irrelevant to your knowing *k*, contrary to what we originally supposed, and there is a threat of contradiction.[24] In the face of this objection, the relevant alternatives theorist may eschew a probabilistic criterion of relevance as such. Yet, it's hard to see what alternative, and otherwise satisfactory, standard of relevance would yield the desired conclusions about the Car Theft Cases, and the value of the relevant alternatives approach in dealing with such cases seems questionable.

An important motivation for pursuing that approach is the hope that this would contribute, down the line, to a solution of the problems raised by Cartesian skepticism. Typically, a relevant alternatives theorist takes the position that we can have knowledge of the external world even though we may be victims of massive sensory deception. On this view, the possibility of such deception leaves our knowledge of the world intact because, with respect to such knowledge, the possibility of deception is an irrelevant alternative. Of course, it won't help just to declare skeptical alternatives irrelevant — that evaluation has to be made in a principled way. Now, suppose that the relevant alternatives approach really did provide an acceptable account of the Car Theft Cases. Such success would mean that relatively pedestrian possibilities like car theft are, in some contexts at least, epistemically irrelevant. All the more reason, then, to hold that the outlandish possibilities raised by skeptics are irrelevant as well.

The envisioned anti–skeptical strategy is to try to assimilate the problem of skepticism to the problem of knowledge in the Car Theft Cases. Such an attempt seems misguided, in light of considerations raised above. The issues arising in the Car Theft Cases have to do with knowledge on the basis of statistical evidence and, perhaps, the requirement of non–arbitrariness in forming justified beliefs. As I have argued, these are not the issues raised by Cartesian skepticism, and there is no reason to expect that a solution to one set of problems will have any bearing on the other set. To be more specific, let's imagine that a preponderance of statistical evidence can create situations in which some alternatives are irrelevant. This is not the situation in which we confront the skeptic (i.e., it's not as though we know, antecedently, that just a handful of the sentient creatures in the universe are massively deceived). So, it isn't easy to see here any basis for the claim that the possibility raised by the skeptic is, for us now, an irrelevant alternative.

VII CONCLUSIONS

I have argued for a number of points concerning the Closure Principle. First, Dretske's Zebra Case does not, on my view, provide a genuine counterexample to the Closure Principle. It seems more plausible that there is a violation of closure in examples like the Car Theft Case. However, even if the Closure Principle does fail in cases of that sort, there is, I maintain, no reason to believe that such a failure carries over to the contexts where the skeptic may appeal to closure. Finally, in my view, serious questions may be raised as to whether the Car Theft Cases really do demonstrate any failure of the Closure Principle at all.[25]

Amherst College

NOTES

[1] This formulation stands in need of further refinements. For, suppose someone knows both *p* and (*p* entails *q*); if that person doesn't put these things together, he or she might fail to infer, and hence not know, *q*. This kind of complication doesn't affect what I want to say below, so I will disregard it. Where a logical consequence is properly recognized as such, I will call it a "clear" logical consequence.

[2] Fred Dretske, [3], p. 1015–1016. Dretske also employs the example in his more recent [5], p. 130.

[3] Dretske, [3], p. 1016.

[4] The problem can't be that you aren't certain that what you see isn't a mule. For, any chance or possibility that the animal is a mule is a chance that it's not a zebra. If this chance makes you uncertain of 'It's not a mule' it should make you equally uncertain of 'It's a zebra'.

[5] Dretske, [3], p. 1016.

[6] Someone might maintain that you don't need this sort of background in formation at the zoo; such information is required out West only because *there* you have information which conflicts with the claim that the animal is a zebra (viz. zebras aren't generally found on Western ranchland). My first response would be that the zoo and ranchland situations are still analogous. If you happen to be at, say, the Bronx Zoo, you have evidence that conflicts with the claim that the animal in the pen is a zebra, namely, the information that zebras aren't native to New York City. In any case, the example could be further modified. Suppose you are in a situation where you mean to identify an animal by sight, but you have no information at all about whether such animals are found in your location, nor about the presence or absence of similar looking but different creatures in the area. Under those circumstances, I think, you couldn't know that the animal is of the sort you would take it to be. I am indebted here to Robert Audi.

[7] For a discussion of these issues, see Robert Stalnaker, [10], especially pp. 63–68.

[8] When I say that there is a statistical reason or statistical evidence in favor of the proposition, I mean roughly the following. Let us say that a statistical probability of an *A*'s being a *B* is one that is assigned on the basis of relative frequencies, counting cases, and so forth. On the basis of such statistical probabilities, a statistical probability may be assigned by direct inference to the proposition 'This *A* is a *B*'. If this statistical probability, in turn, is not zero, we have, other things being equal, some reason—perhaps very small—to think that the *A* in question is a *B*. I am calling such a reason a statistical reason. (My usage here follows John Pollock, [10], p. 231–252).

[9] This analysis will seem misguided to those who doubt that justified acceptance is closed under conjunction. However, it might still be that the existence of the relevantly similar tickets, one of which is known to win, somehow undercuts justification (and knowledge) regardless of how things stand with conjunction. For such a view, see Laurence Bonjour, [1]. The role of the non-arbitrariness constraint in situations like this is also clouded by the fact that someone may fail to know that his or her ticket will lose in lotteries in which the winning chances of the tickets are uneven. I hope to pursue these issues in a further paper; for now, it would be sufficient for my purposes if nothing beyond *statistical probability and abnormality* enters into the proper characterization of these examples. My conclusions below should remain unaffected by dropping any assumptions about the significance of non-arbitrariness in these contexts.

[10] The connection between lottery-like situations and situations where closure (apparently) fails has also been noticed by Jeffrey Olen in [8], p.521–526. I am indebted to David Shatz for this reference.

[11] Compare this set of circumstances with those of a crime-free small town. In a locale where cars are never stolen, you would have no reason at all to think that your car in particular has been stolen, and you can know that it's where you left it. Notice, too, that in such circumstances your car's being taken would be abnormal.

[12] Interestingly enough, the Zebra Case can be made more convincing by filling it out so that a lottery element is introduced. The example could be developed in this way:

Q. Do you know what the animal in the pen is?

A. Sure, it's a zebra.

Q. Do you know for a fact that members of some college fraternity didn't steal the zebra last night as a prank, leaving behind a disguised mule?

The reason one might hesitate to claim to know that such a prank wasn't carried out may be that there is some reason to think that successful, temporarily undetected college pranks are brought off from time to time. Then, in turn, you may not be entitled to say that you know that there isn't a cleverly disguised mule before you. So, it may be that, properly understood or properly filled out, Dretske's Zebra Case should be taken as a member of the family of cases for which the Car Theft Case was the paradigm.

[13] These findings are summarized and discussed by Daniel Kahneman and Amos Tversky, [6].

[14] Which is not to say, of course, that alternative explanations, involving closure failure, can't also be devised. I am indebted here to Richard Feldman.

[15] By a "real" possibility, I mean just one for which there is a positive, even if small, statistical probability; this is a richer notion than plain logical possibility. The associated notion of certainty is the absence of any real possibility of error. This notion of certainty is weaker than the conception of certainty according to which one must have evidence that entails the truth of a belief for that belief to be certain. It is questionable whether the stronger standard of certainty represents a condition for knowledge, since it *ipso facto* rules out the possibility of knowledge by induction. I should make it clear here, though, that I don't intend these glosses to serve as a substantive account of real possibility or of certainty.

[16] The statement of this principle is rough, since it doesn't rule out that the members of L could be entirely unrelated in content. Some stipulation is needed to ensure that L be suitably natural or appropriate; this problem is, of course, closely related to that of choosing an appropriate reference class for direct inference about probabilities.

[17] A similar point is made by Bonjour, [1], p.73n.

[18] Jeffrey Olen suggests that you know the mundane proposition because there is a "nomic connection" between the state of affairs picked out by the propositions which are your evidence and the state of affairs you believe to obtain; in the case of your belief in the clear logical consequence, however, the connection is merely probabilistic and not nomic, and you don't know. Notice, though, that *in the Car Theft Case, it is* nomologically possible for you to have the evidence you have and yet be wrong in your belief about both the initial proposition and the clear logical consequence. So, it is at least obscure exactly how Olen means to draw the crucial distinction. See Olen, [8]. Another explanation of closure failure that would fit the Car Theft Cases is that you "track", in the sense discussed by Nozick, the truth of initial proposition but not that of the clear logical consequence. Nozick's account is presented in his *Philosophical Explanations* (Cambridge: Harvard University Press, 1981); however a discussion of Nozick's work lies outside the scope of this essay.

[19] For more discussion of this possibility, see my doctoral dissertation "Cartesian Skepticism and Epistemic Principles" (Yale University, 1986), Chapter II.

[20] Important early statements of the relevant alternatives theory are found in Fred Dretske, [3] and [4], and in Alvin Goldman [5].

[21] Or, alternatively: the fact that your evidence doesn't exclude the possibility of car theft.

[22] A sophisticated version of this line of thought has been developed by Stewart Cohen in [2].

[23] The relevant-alternatives theorist can't balk at this point, since we're assuming that he or she endorses the Closure Principle.

[24] A related argument may be given to show that the probabilistic criterion of relevance *is unacceptable* when the relevant alternatives theory is couched in terms of reliability over a range of counterfactual situations.

[25] I'm grateful to many people for help in thinking about the issues raised here: Robert Audi, Phillip Bricker, Anthony Brueckner, Fred Dretske, Richard Feldman, John Martin Fischer, Harry Frankfurt, and David Shatz. Recently, I have benefited greatly from conversations with Stewart Cohen.

BIBLIOGRAPHY

[1] Bonjour, Laurence. "The Externalist Conception of Knowledge" in P. French, T. Uehling, and H. Wettstein, eds. *Midwest Studies in Philosophy, Vol. V: Studies in Epistemology* (Minneapolis: University of Minnesota Press, 1980).

[2] Cohen, Stewart. "How To Be A Fallibilist", in J. Tomberlin, ed. *Philosophical Perspectives, Vol. II: Epistemology* (Atascadero: Ridgeview Publishing Company, 1988).

[3] Dretske, Fred. "Epistemic Operators", *The Journal of Philosophy* 69 (1970), p. 1015–1016.

[4] Dretske, Fred. "The Pragmatic Dimension of Knowledge", *Philosophical Studies* 40 (1981), p. 363–368.

[5] Dretske, Fred. *Knowledge and the Flow of Information* (Cambridge: Bradford Books, 1981).

[6] Goldman, Alvin. "Discrimination and Perceptual Knowledge", *The Journal of Philosophy* 73 (1976), p. 771–791..

[7] Kahneman, Daniel and Tversky, Amos. "Choices, Values, and Frames", *American Psychologist* 39 (1984), p. 341–350.

[8] Nozick, Robert. *Philosophical Explanations* (Cambridge: Harvard University Press), 1981.

[9] Olen, Jeffrey. "Knowledge, Probability, and Nomic Connections", *The Southern Journal of Philosophy* 15 (1977), p. 521–526.

[10] Pollock, John. "Epistemology and Probability", *Synthese* 55 (1983), p. 231–252.

[11] Stalnaker, Robert. *Inquiry* (Cambridge: MIT Press, 1984)