Ch006.qxd 2/25/06 9:18 AM Page 83

Chapter 6

Lotteries and the Close Shave Principle

John Collins

1. The Lottery Problem

When I hold a single ticket in a fair lottery with many tickets, then, no matter how probable it is that my ticket will lose, I cannot correctly be said to *know* that it will lose. Call this fact *the lottery observation.* If the moral to be drawn from this observation is that mere probability of truth, no matter how great, is never sufficient for knowledge, then the threat of a new style of skeptical argument looms. For it appears that many of the things I take myself to know are things that I can know only if I can rule out certain quite mundane counter-possibilities, against which my evidence is merely probabilistic.

In recent years this *lottery problem* has assumed a place alongside the Gettier problem and more traditional skeptical considerations (of the Evil Demon or brain-in-a-vat variety) as one of the major challenges facing a satisfactory philosophical account of the concept of knowledge.

In this paper I will argue that the importance of the lottery observation has been exaggerated in the recent literature. The lottery poses no special skeptical threat. For while the lottery observation is quite correct as far as it goes, and illustrates something important about the concept of knowledge, the lottery phenomenon is also a fairly isolated one that fails to generalize in the way that has sometimes been claimed.

I will outline, in Section 2, some of the cases that have been offered as generalizations of the lottery situation. Section 3 presents a characterization of the original lottery case that distinguishes it from the alleged generalizations. I will propose

Aspects of Knowing: epistemological essays
Edited by S. S. Hetherington
Copyright © 2006 by Elsevier Ltd.
All rights of reproduction in any form reserved.
ISBN: 0-08-044979-4

a rather weak necessary condition on knowledge that fails to be satisfied in the lottery example. I claim that a lottery ticket holder fails to know that her ticket will lose precisely because of the failure of this condition. And since this necessary condition on knowing *is* satisfied in each of the supposedly generalized cases, I say that they are not really analogous to the lottery example at all. For all that lottery considerations have shown, it is quite possible that in each of these other examples the subject does possess knowledge. Lotteries pose no special skeptical threat. There is no lottery problem.

This attempt to isolate the lottery example is defended, in Section 4, against an objection stemming from cases of belief based on a certain kind of statistical reasoning.

The aims of the present paper are quite modest. In arguing that there is no special skeptical threat from lottery considerations, I am certainly not suggesting that skepticism is false. And in saying that lottery considerations do not force us to deny knowledge in the allegedly analogous cases, I am not claiming that in all or some of those cases the subject really does know.

2. The Lottery Generalized

It is the belief that the lottery observation readily generalizes that has impressed the lottery problem upon many contemporary epistemologists. Jonathan Vogel was one of the first to suggest that the phenomenon is a widespread one with potentially devastating skeptical consequences. The following example is paraphrased from Vogel (1990).

> *Car Theft*: Suppose that several hours ago Smith left his car parked on a side street in a major metropolitan area. Since Smith clearly remembers where he parked the car, we may be inclined to say he knows where his car is. But does he know that his car has not been stolen in the last couple of hours and driven away from where he parked it? Many people would say that he does not.

Now knowledge is commonly held to satisfy the following closure principle.

> *Closure Principle*: If S knows that p, and S knows that p entails q, then S knows that q.

Given the Closure Principle, there is an obvious tension between the pair of responses Vogel says are standardly elicited by the previous example. That Smith's

Ch006.qxd 2/25/06 9:18 AM Page 85

car is still where he left it entails that it has not been stolen in the last couple of hours and driven away. So from his failure to know that his car has not been stolen and the Closure Principle, we may draw the skeptical conclusion that Smith does *not* in fact know where his car is.

Our primary interest here, however, is with Vogel's reason for thinking that there is an analogy between the Car Theft example and the lottery case that explains why Smith does not know that his car has not been stolen. Here is part of what Vogel has to say on the matter (*ibid*., p. 16):

> 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.

Vogel goes on to offer a range of further examples, which he suggests are also analogous to the lottery case. Here are two of them presented as conversational exchanges (*ibid*., pp. 20–21):

> *Luncheonette*:
> 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*:
> 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.

In a similar vein, John Hawthorne devotes the opening pages of his recent book *Knowledge and Lotteries* to persuading the reader that "the problem posed by lotteries is not an isolated oddity, but is actually widespread" (2004, p. 3, n. 5). The first of Hawthorne's examples approaches the point directly, by drawing attention to the problem of claiming to know *anything* which entails that a particular person fails to win a major prize in a lottery. It should be noted that an example of this sort was already discussed in Harman (1986, pp. 71–72).

Ch006.qxd 2/25/06 9:18 AM Page 86

> *Safari*: Suppose someone of modest means announces that he knows that he will not have enough money to go on an African safari this year. We are inclined to treat such a judgment as true, notwithstanding various far-fetched possibilities in which that person suddenly acquires a great deal of money. We are at some level aware that people of modest means buy lottery tickets from time to time, and very occasionally win. And we are aware that there have been occasions when a person of modest means suddenly inherits a great deal of money from a relative from whom he had no reason to expect a large inheritance. (Hawthorne, 2004, p. 1)

Hawthorne then proceeds to give several other cases that he claims are also analogous to the lottery case, despite the fact that there is no actual lottery involved. I will give only one of these here, to conclude the present section:

> *Heart Attack*: I am [Hawthorne writes] inclined to think that I know that I will be living in Syracuse for part of this summer. But once the question arises, I am not inclined to think that I know whether or not I will be one of the unlucky people who, despite being apparently healthy, suffer a fatal heart attack in the next week. (If only medical self-examination were so easy!) Indeed I am just as unwilling to count myself as knowing about the heart attack as I am to count myself as knowing about the lottery. The analogy continues. Just as I have excellent statistical grounds for supposing that any given lottery ticket will lose, I have excellent statistical grounds for supposing that a given apparently healthy person will not have a heart attack very soon. … And just as many of our ordinary commitments entail that this or that person will lose a lottery, many of our ordinary commitments entail that this or that person will not soon suffer a fatal heart attack. (*Ibid*., 3)

3. The Far-Fetched and the Merely Improbable

Now I think that in fact there is a stark disanalogy between the original lottery case and the attempted generalizations of the previous section. I suggest that this disanalogy is only overlooked when we conflate the notion of a state of affairs being *far-fetched* with the quite distinct notion of its being (merely) *improbable*; when we confuse what is *outlandish* or *far from the truth* with what is just *unlikely*.

Let me introduce the point by examining a case (adapted from Dretske, 1970) that is clearly *not* analogous to the lottery example.

Ch006.qxd 2/25/06 9:18 AM Page 87

> *Zebra*: Smith is at the zoo with his two sons. They are looking at a zebra standing in an enclosure marked "Zebra". Smith's younger son asks, pointing, "Daddy, do you know what kind of animal that is?" Smith replies: "Yes, it's a zebra". Smith's teenage son then turns to his father and says: "But Dad, how do you know that it's not just a mule cleverly painted to look like a zebra?"

My primary concern here is not with whether the teenager's challenge is appropriate, or whether Smith does in fact fail to know that the animal they are viewing is not a painted mule, and hence fails to know that it is a zebra. Neither is the point here to evaluate the status of the Closure Principle, or to ask whether, perhaps, an epistemologically relevant shift of context might have taken place in the course of the conversation. I would simply like to ask the following question: is the Zebra case plausibly construed as a lottery example? In other words: do *lottery considerations* show that Smith fails to know that the animal is not a painted mule?

I think the answer to that question is, quite clearly, no. I know that fraud and deception do take place, even at certain apparently respectable public institutions such as (we may imagine) the zoo in question. But I would deny that, when Smith purchases his admission ticket to the zoo, he is in any sense, in effect purchasing a ticket in a lottery in which some ticket-holders will be picked, essentially at random, to be defrauded and deceived. There is simply no reason to suppose that any such mechanism is in place in the situation described in the example.

The skeptical possibility raised by the teenage son's question is not just improbable, and Smith's evidence against it is not "merely probabilistic" in the same sense in which my evidence that my ticket will lose the lottery is merely probabilistic. Given Smith's background knowledge about the trustworthiness of zoos and of public institutions in general, and of this zoo in particular, the hypothesis is not just very probably false. It is very probably a far-fetched possibility. That is to say, it is a possibility which, very probably, is not *at all* similar to the way things actually are. Smith does not just have probabilistic reasons for thinking that the animal is not a painted mule. He also has probabilistic reasons for believing that the closest alternatives to actuality in which it *is* a painted mule are fairly remote.

Vogel seems to acknowledge this point about the Zebra case. He writes (1990, p. 14):

> 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

Ch006.qxd 2/25/06 9:18 AM Page 88

> 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.

But now contrast this with the Lottery case. While the possibility that my single ticket will be the winner is very improbable, it is also a possibility that I know to be very close to actuality (if not actual!) and not at all far-fetched. At least, that is the case if I know that the lottery is a fair one.

I propose that the reason I fail to know that my ticket will lose is that I know that, no matter what happens, I will either win, or, at worst, come very *close* to having won.

In other words, I am suggesting that the reason I do not know that my ticket will lose, despite the overwhelming probability I assign to that prospect, is that the following necessary condition on knowing is not satisfied:

> *The Close Shave Principle*: If S knows that p, then there is no possibility that is very close to actuality at which p is false and to which S assigns non-zero probability.

Equivalently:

> If there is some possibility that is very close to actuality at which p is false, and to which S assigns non-zero probability, then no matter how subjectively improbable this possibility is, S doesn't know that p.

Now let us return to the first of the supposed analogies to the Lottery: Vogel's Car Theft case. If my suggestion is correct, then much will depend on how we fill in the details of the example.

First of all, suppose that, in the neighborhood in which Smith parked his vehicle, there is a gang that, quite literally, steals cars at random. Each evening they commence by drawing a ball from an urn to determine which side street they will hit. Then they toss a coin to decide whether to take a car from the north or south side of the street, then another ball is drawn to determine which car on that side of that street will be taken. If this is the case, then, yes, if Smith parked in that neighborhood, he certainly does not know that his car has not been stolen, and for the same reason that I do not know that my ticket will not win the lottery.

Ch006.qxd 2/25/06 9:18 AM Page 89

Alternatively, let us suppose that there is indeed a car thief at work in the neighborhood, that she operates opportunistically, if not completely randomly, and that as a matter of fact she comes very close indeed to stealing Smith's car on the night in question. Walking along the street, trying car door handles, ignoring those vehicles with obvious anti-burglary devices, she comes upon Smith's car. It is locked but unprotected. She sees that the coast is clear, but then, just as she is about to force the lock, a police patrol car swings around the next corner. With the police car's headlights upon her, she casually resumes her stroll down the block. By the time danger has passed, she has happened upon another vehicle just as attractive as Smith's. She steals that car instead.

In this last scenario, does Smith, sitting blithely at a nearby restaurant, know that his car has not been stolen? I am strongly inclined to say: "No, he doesn't't". This example provides both illustration and confirmation of the Close Shave Principle, I think. Were an observer of the whole episode later to describe to Smith exactly what had transpired, Smith might well respond as follows: "Whew! That *was* a close thing. I thought I knew where my car was that night when I left the restaurant, but I guess I did not. I was just lucky that it was still there where I parked it".

Absent such details, what should we say about the example though? I suppose I am inclined to agree that Smith does not know that his car has not been stolen, but the situation is really not at all clear to me, and neither is it my present concern to settle the matter. What is clear to me, however, is that lottery considerations alone do not force us to say that Smith lacks knowledge in the unadorned Car Theft case. If we do not *know* that Smith's car comes close to being stolen, then we cannot appeal to an alleged analogy with the Lottery case in order to deny Smith knowledge of his car's whereabouts.

What of the more direct kind of example that Hawthorne adapts from Harman, in which a claim has been made to know something that entails that a particular person will not be a lottery winner?

Consider the Safari example. Suppose that Smith, being of modest means, announces that he *knows* he will not be able to afford an African safari this year. Should we concur with this judgment? Note that in order to cast skeptical doubt on Smith's claim via lottery-style considerations it is not sufficient simply to remind ourselves (as Hawthorne does) that "people of modest means buy lottery tickets from time to time, and very occasionally win". We would need to know in addition, for example, that *Smith* is in the habit of buying the occasional lottery ticket, or that someone will give (or very nearly might have given) Smith a lottery ticket in the near future. After all, if I know that someone *does not* hold a ticket in a certain lottery, then there is no apparent problem with my knowing that he will not win that lottery.

Ch006.qxd 2/25/06 9:18 AM Page 90

In short, if I know that Smith does have a lottery ticket, then we are back in the original lottery case. We should then simply agree that Smith does not know that he will not win a major prize, and hence that he does not really know that he will be unable to afford the safari. If, on the other hand, the story leaves it unclear whether or not Smith will have, or very nearly might have had, a lottery ticket, then Smith's claim to knowledge is as unclear as in some of the previous cases. But then the example is clearly no longer one to which the lottery observation straightforwardly applies.

Similar comments apply to the Heart Attack example. Heart attacks are not the result of any process that remotely resembles an actual lottery. If I am generally in good health then I consider it not only very probable that I will not suffer a heart attack in the next few months, but also very probable that I will not even come close to suffering a heart attack. If I later learn than I *did* come close, perhaps, say, because a crucial artery was very nearly blocked by a blood clot, then I will certainly retract any claim to have known that I would not. Once again, this serves to confirm the significance of the Close Shave Principle.

Of course, this is not to say that in the absence of any particular medical evidence it *is* ever correct to say that I know that I will not suffer a heart attack in the near future. All I am claiming here is that lottery considerations are not relevant to the example. I suspect that our reluctance to ascribe knowledge in this sort of case, and to seek a medical check-up even when in apparent health, has more to do with the high cost of being wrong, than to any analogy with the case of the lottery.

4. Inductive Knowledge

Suppose that on the basis of empirical observation we have come to believe that a certain generalization is both true and law-like. It appears that such a belief, acquired by inductive inference from a limited sample, can straightforwardly count as an item of knowledge, despite the fact that what we know about the sample fails to entail the truth of the generalization. And in many cases the possibility of acquiring such knowledge is quite unproblematic for the view being defended here. If, for example, we have come to believe on the basis of what has been observed, that, as a matter of physical law, metals are good conductors of electricity, we will not only assign a high probability to the proposition that this (untested) piece of metal will prove to be a good conductor; we will also consider it highly probable that the closest alternatives to actuality at which the metal strip is *not* a good conductor will be rather remote possibilities.

An apparent problem arises, however, in the case of law-like statements arrived at inductively, that turn out on closer examination to be statistical in nature.

Ch006.qxd 2/25/06 9:18 AM Page 91

Consider for example the following two cases:

> *Ice Cube*: Several hours ago, you left an ice cube in a thermos half-filled with lukewarm water. It occurs to you that the ice cube you left in the lukewarm water must have melted by now. Despite the fact that you are not presently looking at the contents of the thermos, you know that the ice cube has melted. (cf. Vogel, 1987, p. 206)
>
> *Uranium*: Roger places a piece of uranium on a photographic plate, and discovers that the plate has become fogged. He repeats the experiment many times. After numerous trials, he puts a piece of uranium on a plate, goes away from his laboratory, and returns some time later. Roger believes that the plate is fogged. Moreover he knows, by induction, that the plate is fogged, even before he inspects it. (see Vogel's manuscript)

The problem here is that physicists tell us that in each of these examples, the process that underwrites the inductive generalization (ice melts in the sun on hot days, uranium fogs photographic plates) is a merely statistical one. Does it follow from the Close Shave Principle that neither is a case of knowledge? If so, then the Close Shave Principle stands refuted.

Now in order to see whether the Principle really applies to cases such as these, we will have to look more closely at the particular details of the examples.

Let us take the uranium example first. Suppose that Roger placed a one-gram lump of uranium on a photographic plate several hours ago. The dominant isotope that makes up more than 99 percent of the element as it occurs naturally is Uranium 238, which has a half-life of about four and a half billion years. The "half-life" is a period of time in which half of the atoms present may be expected to have undergone radioactive decay. Because the half-life in this instance is so long, the probability that any single particular atom in the lump will have decayed in the few hours that Roger has left it unattended is extremely small. And, since the decay process is an essentially indeterministic one, this is the kind of situation to which the Close Shave Principle has application.

Suppose Roger is somehow able to form the belief, *de re*, about some single particular atom in the lump, that in the several hours since he left the piece of uranium unattended in his lab, that atom has not decayed. Then this is indeed belief in a "lottery proposition" (in Vogel's terminology). Roger's belief is very probably true, yet, even supposing that the belief actually *is* true, Roger does not *know* that the atom has not decayed in the last several hours. Why? Because to suppose that, contrary-to-fact, the atom in question had decayed, requires us only to envisage

Ch006.qxd 2/25/06 9:18 AM Page 92

a possible state of affairs that differs from actuality with respect to one tiny particular matter of fact. This is, then, a possibility that is very close to actuality, to which Roger assigns a (very small yet) positive probability, and in which Roger's belief about the atom is false. The Close Shave Principle tells us that no matter how improbable Roger considers this possibility to be, he does not know that it does not obtain.

But Vogel's example is quite different in structure to that. It involves the indeterministic behavior over several hours, not of a single atom, but of a vast quantity of atoms. There are so many atoms in a one-gram lump of uranium that, even given the very small probability that any particular atom will undergo decay, one may confidently expect about 25,000 of these individually improbable events to occur each second that passes. Each of these decay events involves the emission of an alpha particle accompanied by weak gamma radiation. Over the course of several hours there will be hundreds of millions of individual decay events within the lump, and this suffices to expose the photographic film.

So, in order for the photographic plate not to have become fogged since Roger left it, some hundreds of millions of independent chance events that actually did take place must be supposed not to have taken place. But this means that there is no possibility that is close to actuality in which the photographic plate is not fogged, since any possibility that differs from actuality with respect to the occurrence of some hundreds of millions of independent events is not at all close to the way things actually are. The proposition that the plate has not become fogged is not a "lottery proposition" at all; the Close Shave Principle simply does not apply here.

Consistently with our observations about the subject's lack of knowledge in the original lottery case, we can say the following things about this example: Roger knows that it is physically possible, though overwhelmingly improbable, that the photographic plate is unfogged. Roger knows, moreover, that it is overwhelmingly probable that the plate is *not even close* to being unfogged, i.e. that there is no possible world at all close to the actual world at which the plate is not fogged. So, consistently with the Close Shave Principle, we say in this case that Roger does know that the plate has been fogged by the uranium.

Similar considerations apply to the case of the ice cube. It is true that there are physically possible initial microstates of the system consisting of the water and the ice that are (i) macroscopically indistinguishable from the actual initial microstate of the system, and which (ii) evolve according to the laws of physics over the next several hours into a state in which the ice remains unmelted in the water. (By "microstate" here I mean a complete specification of the initial position and momentum of each and every one of the water molecules in the thermos.) Do I know that the system was not initially in one of those unusual microstates?

Ch006.qxd 2/25/06 9:18 AM Page 93

It should be clear by now that lottery considerations do not force a negative answer to that question. For it is not just that I consider the possibility that the ice remains unmelted overwhelmingly improbable. In addition I think it overwhelmingly improbable that this possibility is even remotely like the way things actually are in the thermos. There is no small adjustment to the actual initial state of the system that would transform it into one in which the ice fails to have melted after floating for several hours in the lukewarm water. That would require a quite extraordinary coordination of the trajectories of the individual particles in the system. Such extraordinary coordination is physically possible, but this possibility differs from the way things actually are with respect to the positions and velocities of a vast number of molecules. Once again: any possibility that differs from actuality with respect to a vast number of independent matters of particular fact is not a possibility that is even remotely close to the way things actually are. The proposition that the ice remains unmelted might be assigned some tiny positive probability, but it is not a "lottery proposition".

I have argued that lottery considerations pose no threat to this sort of inductive knowledge. But is the discussion of the preceding paragraphs consistent with my concession that in the original, genuine, lottery example, the ticket-holder fails to know he will lose? In other words: can the original lottery example survive the same level of scrutiny to which the Uranium and the Ice Cube examples have just been subjected?

Consider a typical lottery mechanism. In the New York State Lotto, for example, ping pong balls bearing numerals from "1" to "59" are released into a transparent chamber. Jets of air are blown into the chamber to mix the balls. A valve is then opened, through which six balls are allowed to pass into a clear tube that leads to the display area.

The air-mixing mechanism is of crucial importance. The resulting jostling of balls is a dynamic process, which though quite deterministic, is also highly modally sensitive. Very slight changes to the positions or velocities of just a few balls will be amplified by the ensuing collisions to produce completely different results. This ensures that every possible combination of numbers will, at worst, have come very close to having won, no matter what actually transpires when the valve is opened. One consequence of this is that predicting the result of the draw in advance is a practical impossibility. Another consequence is that, by the Close Shave Principle, no ticket holder can know in advance of the draw that her ticket is going to lose. As Hawthorne (2004, p. 8) reminds us, the official slogan of the New York State Lottery is: "Hey, you never know".

Note that it would be not at all appropriate here to argue that making, e.g. a slight change to the position of single ball would really require changing a vast number of particular matters of fact, since each ball is itself made up of a vast quantity of particles. The ball is a rigid body whose parts move together in a quite predictable way. The particular matters of fact about the precise locations of the micro-particles that make up the ball are not *independent* matters of particular fact.

Another example discussed by Vogel (manuscript) seems to be poised delicately between the genuine Lottery case and examples of the statistical-inductive kind just considered:

> *Hole-in-One*: Sixty golfers are entered in the Wealth and Privilege Invitational Tournament. The course has a short but difficult hole, known as the "Heartbreaker". Before the round begins, you think to yourself that, surely, not all sixty players will get a hole-in-one on the "Heartbreaker".

Do you know that not all the players will get a hole-in-one? Vogel thinks that you do.

I am inclined to agree with Vogel here, but for what I suspect are rather different reasons. In fact I believe that the example is under-described. I think that judgment as to the correctness of an ascription of knowledge might well go either way in this case, depending on how various missing details of the story are filled in.

Let us suppose that the players in the tournament are professionals. When a skilled professional golfer aims to land his ball on the green on a par three hole, then whether or not he succeeds is typically not at all a matter of chance. In fact a professional golfer may well be able to land the ball reliably not only on the green, but on a certain part of the green. Often a player will aim his shot right at the pin; that is, attempt to "hole out", or, failing that, to leave the ball as close to the hole as possible. Now of course not even the best golfers have the level of control that would be required to hole out reliably at a distance of, say, 100 m. So we might well take it to be a matter of chance whether or not a certain skilled golfer who is attempting to make a hole-in-one on the Heartbreaker will be successful or not. The proposition that this player makes a hole-in-one is indeed then a lottery proposition, and I do not know in advance that she will not succeed. I think she will succeed in landing the ball on the green, perhaps quite close to the hole. I doubt very much that the ball will actually land *in* the hole, but I do not know that it will not. That is because, although I doubt very much that the shot will result in a hole-in-one, I am also fairly sure that there is a possible state of affairs that is very close to the way things actually are, one perhaps in which the trajectory of her swing is ever-so-slightly different, in which the player *does* make a hole-in-one.

Ch006.qxd 2/25/06 9:18 AM Page 95

Matters are complicated, however, by the fact that, not infrequently, a good golfer will not even be attempting to make a hole-in-one when she steps up to the tee at a par three hole. Suppose, for example, that the pin is positioned near an edge of the green that is bordered by a deep and treacherous sand trap. Then it may make more sense strategically to aim for a safer region of the green, from which position the player can be confident of completing the hole in at most two putts, rather than to run the risk of disaster. If that is what our golfer has in mind when she plays the tee shot, then I think our epistemic judgment about the example changes. She is aiming the ball at an area of the green some distance from the pin, and, since she is a skilled player, it is very probable that she will succeed in landing the ball somewhere in that region. We can then be confident that she will not make a hole-in-one, because doing so would now require her to *mis-hit* the ball, something that she will probably not even come close to doing. In this situation, the proposition that the player does not make a hole-in-one is no longer a lottery proposition at all, and lottery considerations give us no reason at all to think that one cannot know that she will not make a hole-in-one.

Returning to Vogel's own version of the example: the reason I think you do know that not all 60 players will make a hole-in-one on the "Heartbreaker" is that it is highly probable that among those 60 players in the tournament, several (at least) will, for one reason or another, be trying *not* to make a hole-in-one, and can be reliably counted on not to make a hole-in-one when attempting not to do so. These players might be attempting, rather, to place the ball in a position from which the hole can be completed safely in par. The particular details of the example-scheme might be filled out in many ways, of course, but it seems to me that for many of the most natural looking scenarios, Vogel is mistaken in thinking that the case is "very much like the lottery".

5. Conclusion

In the preceding sections various examples have been examined in considerably more detail than is customary, but if the central thesis of this paper is correct, then it is precisely those details that are going to determine, in any particular case, whether or not the subject can be correctly said to know.

The suggestion that lottery considerations provide an alternative route to skepticism depends on the claim that lottery propositions are widespread. I doubt very much that that is so. Such propositions will be found only where there is an underlying mechanism that ensures the existence of very close alternatives to actuality in which the lottery proposition is false. In many of the ordinary examples that have been offered as generalizations of a genuine lottery situation, it is quite

Ch006.qxd 2/25/06 9:18 AM Page 96

implausible to imagine that there is any such underlying mechanism. While it is true that a ticket-holder in a fair lottery does not know that his ticket will lose, the epistemological significance of this observation is, it seems to me, rather slight.

References

Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy*, *67*, 1007–1023.

Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Clarendon Press.

Vogel, J. (1987). Tracking, closure, and inductive knowledge. In: S. Luper-Foy (Ed.), *The possibility of knowledge: Nozick and his critics* (pp. 197–215). Totowa, NJ: Rowman & Littlefield.

Vogel, J. (1990). Are there counterexamples to the closure principle? In: M. Roth & G. Ross (Eds), *Doubting: Contemporary perspectives on skepticism* (pp. 13–27). Dordrecht: Kluwer.

Vogel, J. (manuscript). Subjunctivitis. AQ1