

KNOWLEDGE AND LOTTERIES

DONALD SMITH

Virginia Commonwealth University

John Hawthorne's recent monograph *Knowledge and Lotteries*¹ is centred on the following puzzle: Suppose you claim to know that you will not be able to afford to summer in the Hamptons next year. Aware of your modest means, we believe you. But suppose you also claim to know that a ticket you recently purchased in a multi-million dollar lottery is a loser. Most of us have the intuition that you *do not know* that your ticket is a loser. However, your alleged knowledge of not being able to afford to summer in the Hamptons puts you in a position to know that your ticket is a loser. For the proposition that you will not be able to afford to summer in the Hamptons entails the proposition that you will lose the lottery. And the following principle, what Hawthorne calls 'Single Premise Closure' (p. 34), is very plausible: If you know that p , p entails q , and you competently deduce q from p thereby coming to believe that q (all the while retaining your knowledge of p), then you come to know q .

Perhaps you think the lottery case is not so very puzzling. Shouldn't we simply retract our belief in your claim to know that you will not be able to summer in the Hamptons? That alone would not commit us to anything objectionable such as scepticism about most of what we ordinarily take ourselves to know. However, as Hawthorne notes—and as other philosophers have noted, as well²—the lottery case appears to objectionably generalise. For instance, consider the following cases.

The Heart Attack Case. You claim to know that you will dine with a friend next month. But it seems you *don't know* that you will not be one of the unfortunate apparently healthy people to die of a sudden heart attack before then. (Along the same lines, it seems that you *don't know* that you will not be involved in a fatal car accident before your scheduled dinner, that you *don't know* that you will not suffer a fatal and unexpected brain aneurysm before then, etc.) The proposition that you will have lunch with a friend next month entails that you will not die of a sudden heart attack before then. And once again by Single Premise Closure, you can come to know that you will not be one of the apparently healthy people to die of a sudden heart attack before next month. In this case, denying the original knowledge claim may seem less appealing.

The Quantum Case. You claim to know, by present perception say, that there is a table in your office. According to quantum mechanics, there is a wave function that describes the nomically possible developments of those particles

1. John Hawthorne, *Knowledge and Lotteries* (Oxford University Press, 2004. x + 206 pp. £25.00).

2. See, for instance, Stewart Cohen, 'Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery', *Australasian Journal of Philosophy*, 76 (1998), pp. 289–306, and Jonathan Vogel, 'Are there Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.), *Doubling: Contemporary Perspectives on Scepticism* (Kluwer, 1990).

that compose the table. On standard probabilistic interpretations of quantum mechanics, there is a non-zero probability that the particles that compose the surface of the table remain in more or less the same arrangement while the remaining particles that compose the inside of the table become arranged in such a bizarre way that the particles no longer compose a table. Say that such particles in such a situation compose a table façade. Now, it seems that you *don't know by present perception* that this is not one of those quantum occasions in which some particles have come to compose a table façade. But, like the previous cases, given your alleged perceptual knowledge and Single Premise Closure, you can come to know that this is not such an occasion. Denying the original knowledge claim in this case is very unappealing. Doing so would amount to embracing scepticism about all our beliefs based on present perception.

Knowledge and Lotteries, then, is focused on a genuine epistemological puzzle. Chapter 1 provides an introduction to the above cases, as well as others, and the general structure of the puzzles: In each case, there is an ordinary proposition that we typically take ourselves to know, a lottery proposition that—though in some sense probable—we do not take ourselves to know, the ordinary proposition entails the lottery proposition, and Single Premise Closure puts us in a position to know the lottery proposition.

Hawthorne also undertakes three other important tasks in the opening chapter. First, he offers a diagnosis of our intuition that we do not know lottery propositions. Hawthorne settles on an account in terms of parity reasoning. Roughly, lottery scenarios affect a conceptualisation of possible outcomes where one has no significantly better evidence for thinking that one of the outcomes will/will not obtain than one has for thinking any other outcome will/will not obtain.

Second, Hawthorne traces connections between knowledge, assertion, epistemic modality, and practical reasoning in terms of principles that serve as desiderata for solutions to the lottery puzzles. Roughly, the principles are: S ought to assert that p only if S knows that p , it is epistemically possible for S that p only if p is compatible with what S knows, and S ought to use only what S knows in S's practical deliberations.³

Third, Hawthorne introduces Single Premise Closure, ably defends it against notable objections from Dretske⁴ and Nozick⁵ and provides compelling positive reasons on its behalf. Another closure principle, Multi-Premise Closure, is also introduced and defended. According to Multi-Premise Closure, if S knows p_1, \dots, p_n , competently deduces q thereby coming to believe q while retaining knowledge of p_1, \dots, p_n , then S comes to know q . Consistency with these closure principles—especially Single Premise Closure—also serves as a desideratum for solutions to the lottery puzzles.

3. Hawthorne discusses other desiderata for solutions to the lottery puzzles, for instance, respecting a principle tying epistemic probability to objective chance, van Frassen's so-called Reflection Principle, and a disquotation schema for the predicate 'know(s)'.
4. 'Epistemic Operators', in *Perception, Knowledge and Belief* (Cambridge University Press, 2000).
5. *Philosophical Explanations* (Oxford University Press, 1981).

The remainder of *Knowledge and Lotteries* is devoted to assessing various solutions to the lottery puzzles especially with regard to whether the aforementioned desiderata are respected. In Chapter 2, Hawthorne considers contextualist solutions. According to the contextualist, there is no unique semantic value for every token of the predicate 'know(s)'. The semantic value for tokens of 'know(s)' varies from context to context in such a way that different utterances of a sentence of the form 'He knows that *p*' can have different truth values even if the occurrence of 'He' has the same referent. According to the contextualist, our apparently incompatible intuitions in lottery cases are to be explained by subtle shifts in certain contextual parameters that result in semantic shifts for 'know(s)'. So, for instance, in 'ordinary' contexts we regard an utterance of 'You know that you will dine with a friend next month' as true. In another context, we regard an utterance of 'You know that you will not be the victim of a sudden heart attack before next month' as false. These different contexts affect a shift in semantic value for the different occurrences of 'know' in the respective utterances. Roughly, in the first context, a less strict or 'low standards' semantic value is expressed by 'know' and in the second, a more exacting or 'high standards' semantic value is expressed.

Much of Chapter 2 is devoted to a powerful litany of criticisms of contextualism. Of particular interest—especially to contextualists—should be Hawthorne's novel objection that adopting contextualism about 'know(s)' puts considerable pressure upon one to adopt contextualism about 'ought to assert' and 'ought to use as a premise in practical deliberation'. According to Hawthorne, however, these normative predicates do not seem context sensitive. Put differently, what is proper to assert and what is proper to use as a premise in one's practical deliberations do not seem to vary from context to context. Contextualists will no doubt attempt to find some wiggle room. Nevertheless, the objection is an especially important one since contextualists have appealed to an intuitive tie between knowledge and assertion to motivate their view.⁶

Two invariantist solutions—viz. sceptical invariantism and simple moderate invariantism—are discussed in Chapter 3. Roughly, according to invariantism, 'know(s)'—at least as used in standard 'knows that' constructions—has a single semantic value that does not vary from context to context. According to sceptical invariantism, most, if not all, of our claims to know are false. So, the sceptical invariantist replies to the lottery puzzles by denying knowledge of the ordinary propositions, for instance by denying that you know that you will dine with a friend next week and that you know that you are looking at a table. Sceptical invariantism does poorly on Hawthorne's score card: it denies most of what we ordinarily take ourselves to know and it disrupts the conceptual connections between knowledge, assertion, and practical reasoning; if we do not know most of what we claim to know, then—given the principles tying knowledge to assertion and practical reasoning—most of our assertions

6. See, for example, Keith DeRose, 'Assertion, Knowledge, and Context', *Philosophical Review*, 111 (2002), pp. 167–203.

are unacceptable and should not be employed as premises in our practical deliberations.

According to simple moderate invariantism, most ordinary knowledge claims are true, but we know or at least can come to know lottery propositions. This version of invariantism also fares poorly with respect to assertion and practical reasoning. If you know that your lottery ticket will lose, then why does it seem inappropriate for you to assert that you will lose? To borrow an example from Hawthorne (p. 29), suppose you are offered a penny for your lottery ticket and you reason as follows:

The ticket is a loser.
So if I keep the ticket I will get nothing.
But if I sell the ticket I will get a penny.
So, I should sell the ticket.

This seems an obviously bad piece of practical reasoning. But the simple moderate invariantist is hard pressed to explain why the above reasoning is fallacious. For he thinks that you do know or can come to know that your ticket is a loser. Simple moderate invariantism also appears to violate Multi-Premise Closure. If you can know that your ticket will lose, then in principle, it seems that you can come to know for any particular ticket that it will lose. And by Multi-Premise Closure you will then be able to come to know that every ticket will lose, which is unacceptable.

Finally, in Chapter 4, Hawthorne develops an alternative version of invariantism—sensitive moderate invariantism—and an alternative solution to the lottery puzzles. According to the contextualist, contextual shifts result in semantic shifts. According to the sensitive moderate invariantist, varying contextual parameters do not result in semantic shifts for ‘know(s)’ but do affect whether a subject knows that *p*. Much of Chapter 4 is spent developing two such parameters, viz., salience of counterpossibilities and practical environment. I focus here on practical environment because according to Hawthorne (pp. 173–174) salience does not provide the key to solving the lottery puzzles. (That said Hawthorne has much of interest to say about when and whether certain counterpossibilities are salient.) Here is Hawthorne on the mechanism of practical environment and how it is relevant to whether a subject knows that *p*:

Insofar as it is unacceptable—and not merely because the content of the belief is irrelevant to the issues at hand—to use a belief that *p* as a premise in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time. Thus when offered a penny for my lottery ticket, it would be unacceptable to use the premise that I will lose the lottery as my grounds for making such a sale. So on that occasion I do not know that I will lose. (p. 176)

As I understand him, Hawthorne is offering a necessary condition for a subject to know that *p* on a particular occasion, something like the following:

The Practical Reasoning Constraint (PRC): Suppose that p is practically relevant for S at some time t . Then, if it is unacceptable for S to use the premise that p in a piece of practical reasoning at t , then S does not know that p at t .

Furthermore, as I understand Hawthorne, he is considering the suggestion that PRC provides the key to solving the lottery puzzles.⁷ So Hawthorne:

We now have the beginnings of a diagnosis of our epistemological puzzle-ment. We underestimate the contribution of practical environment to the truth of knowledge ascriptions. The picture just given is compatible with the idea that most ordinary knowledge claims come out true. But when we reflect as philosophers, it does not occur to us that issues about practical environment may be relevant to the truth of those ordinary ascriptions. We are insensitive, and attempt to evaluate knowledge ascriptions out of context. In particular, we fail to consider the deliberative context of the subject. No wonder we get confused. (p. 180)

It is not clear to me, though, that PRC unlocks the door to our epistemological puzzlement; indeed, I shall argue that it does not. For concreteness, let us focus on one such puzzle, the Heart Attack Case. What exactly is the solution being proposed to this puzzle? We can perhaps make some headway towards an answer by supposing that someone offered you life insurance and you reasoned as follows:

I shall be dining with a friend next month.
So I shall not die before then.
So I ought to wait to buy life insurance at least until then.

This too is a bad piece of practical reasoning. Given our supposition, whether you will be dining with a friend next month is practically relevant and it is unacceptable for you to use the premise that you will be dining with a friend next month in the above piece of practical reasoning. So, by PRC, it follows that you do not on that occasion know that you will be dining with a friend next month.

The above remarks are all well and good, but they fail to indicate what goes wrong in the Heart Attack Case. This is so for the simple reason that in the case as described whether you will be dining with a friend next month is not practically relevant, at least not in the sense important for Hawthorne's proposal. In the case as described, you are not offered life insurance and you do not engage in any practical deliberations. You claim to know that you will be dining with a friend next month. We point out to you that if that were true, you would be in a position to know that you will not suffer an unexpected and fatal heart attack before then. But, we continue, it seems implausible to

7. Hawthorne never explicitly endorses this solution to the lottery puzzles, but he clearly prefers it. See pp. 187–188.

say you are in a position to know the latter and conclude that you do not know the former. As far as I can see, in the Heart Attack Case, PRC is unable to get a purchase because practical environment does not play a relevant role in that case. Moreover, the original lottery case and the Quantum Case are, in that respect, like the Heart Attack Case. You do not engage in any practical reasoning in the Quantum Case and there is nothing analogous to someone offering you life insurance or a penny for your lottery ticket. PRC, then, is silent about what to say in the Quantum Case as well.

This problem is not a trifling one. The only way I can see of getting around it is to drop or weaken the antecedent condition of PRC having to do with whether or not p is practically relevant for S at t . Dropping said condition, however, takes the moderation out of the proposed solution. For presumably, the above pieces of fallacious practical reasoning are not unacceptable on some occasions and acceptable on others; rather, they are fallacious *simpliciter*. But then by the principle that results from dropping the aforementioned antecedent condition of PRC, it follows that you do not know *at any time* that you will be dining with a friend next month. Furthermore, I do not see any way of weakening the condition that avoids such scepticism. A counterfactual understanding of whether or not p is practically relevant for S at t will not help. On an occasion where the moderate invariantist wants to say that you do know that you will dine with a friend next month, it will still be true that *were* you offered life insurance and reasoned in the above way, it would be unacceptable.

I cannot see how to amend PRC in such a way that it proves helpful in providing a *moderate* solution to the lottery puzzles with which we began. Perhaps this is a case of philosophical short-sightedness on my part. At the very least, however, more needs to be said about how exactly practical environment resolves the epistemological puzzlement induced by the lottery puzzles in question.

Turning toward another criticism, I would like to have seen more discussion of whether the original lottery case *objectionably* generalises. In particular, I would like to have seen more discussion of the generalisation to the Quantum Case. Scepticism about most of my beliefs about the future does not strike me as too implausible. Indeed, it seems to me fairly plausible to say that I do not know that I shall be alive tomorrow given all the myriad live possibilities the obtaining of which entail my death. Such possibilities are, hopefully, unlikely and so for practical purposes I can ignore them. Nevertheless, it seems right to say that I do not know whether I shall be alive tomorrow. I would be much less willing, however, to live with scepticism about all my beliefs based on present perception. I am unconvinced, however, that the Quantum Case calls my beliefs based on present perception into question.

Suppose that a moment ago some particles in my office in fact composed a table but have since then undergone the sort of bizarre rearrangement described in the Quantum Case. It is not at all clear that I no longer perceive a table. We are told in the Quantum Case that such a system of particles is not a table. But why think that? The present relevant system of particles in my office is not anything like a hologram or a hallucination or an appearance

caused by a malicious neuroscientist. Our concept of a table is no doubt complicated and I doubt whether anyone has ever given informative criteria for being a table. Nevertheless, our concept of a table is clearly a functional one and I see no reason for thinking that an object composed by particles arranged in the way we are presently imagining could not fulfil the relevant functions. For instance: a moment ago, my computer, printer, coffee mug, and feet were being supported by a table. There is no reason to think that after the bizarre rearrangement of particles, those objects would fail to be so supported. I see no reason, then, for thinking that such a system of particles would not satisfy our concept of a table. And so, I see no reason for thinking that the Quantum Case jeopardises my particular claim to know that there is a table in my office or putative perceptual knowledge in general.⁸

Set aside the above remarks. An alternative reply is to try to point out a relevant difference between standard lottery cases and the Quantum Case. Earlier I mentioned Hawthorne's parity reasoning proposal about our inclination to deny knowledge of certain lottery propositions. In lottery cases, such parity reasoning is buttressed by the justified belief that one of the possible outcomes will obtain, for example that it is buttressed by our justified belief that one of the tickets will be drawn. (Similarly, in the Heart Attack Case, parity reasoning is supported by our justified belief that some apparently healthy person will die of a sudden and fatal heart attack.) Because I am justified in believing that some ticket will win and I know that I have a ticket, I recognise that there is some reason for thinking that my ticket will win. But I haven't any better evidence for thinking that my ticket will not win than I do for thinking that any other ticket will not win. So, I am not able to know that my ticket will not win.

Here I think the Quantum Case diverges from the other lottery cases. The standard interpretation of quantum mechanics does not provide us with a reason for thinking that some system of particles has been or is or will be a table façade. It is a theoretical prediction of this interpretation of quantum mechanics that the wave functions that describe the nomically possible developments of physical systems include a very small proportion of bizarre developments of those systems. This theoretical prediction is the basis of the non-zero (but exceedingly small) probability of a physical system being a table façade or a desk façade or a tree façade, etc. We do not, however, on the basis of this theoretical prediction alone have a good reason for believing that in the past there has been or at the present moment there is or in the near future there will be a table façade or a desk façade or any other kind of physical system façade. If physicists were to empirically confirm the existence of a table façade, then we may have some reason for thinking that there are or will be

8. Perhaps the Quantum Case raises questions about my belief—supposing I have this belief—that the table I am presently perceiving is the *same* table I perceived a moment ago. Perhaps the table in my office a moment ago could not survive the bizarre rearrangement of particles that once composed it. But this too is not obvious to me. At any rate, a discussion of what sorts of changes material objects can and cannot survive, while metaphysically interesting, is beyond the scope of this paper.

other such physical systems. But, as far as I am aware, there has not been an empirically confirmed instance of such a physical system façade. I do not think, then, that the kind of parity reasoning used in the lottery cases, gets a foothold in the Quantum Case. I have very good perceptual evidence that there is a table in my office and no good reason for thinking that some alleged and otherwise apparent table is not in fact a table but a table façade. Accordingly, I have significantly better evidence for thinking that there is a table in my office than I do for thinking that there is a mere table façade in my office.

It will not do to emphasise that wave functions of physical systems describe their *nomically* possible developments. For other more traditional sceptical hypotheses such as brain-in-a-vat hypotheses are thought to be nomically possible, as well. But as Hawthorne notes (p. 6) many think we have significantly better evidence for thinking that I am perceiving a table than that I am a brain in a vat merely seeming to perceive a table. This brings me to my final criticism. Hawthorne suggests that the lottery puzzles constitute a different and more challenging obstacle to knowledge claims than the one posed by more traditional sceptical arguments. He says:

Indeed, lottery-style considerations are arguably a more dialectically effective tool for the skeptic than standard brain in a vat or deceiving demon thought experiments. While many contemporary philosophers are inclined to resist the skeptic by claiming that they can, after all, know that they are not brains in vats, they are not nearly so eager to embrace the claim that they know they will lose a lottery for which they hold a lottery ticket. And once this has been conceded, it is extremely hard to justify a different attitude to the other ‘lottery propositions’ that figure in the above examples. (p. 6)

But is there any epistemically relevant difference between a brain-in-a-vat sceptical argument and Hawthorne’s Quantum Case? The above remark about the nomological possibility of my being a brain in a vat suggests that the answer is ‘no’. For given the nomological possibility that I am a brain in a vat, there is some non-zero objective probability that I am a brain in a vat. But it is exceedingly difficult to make out any epistemically relevant difference between there being a non-zero objective probability that I am perceiving a table façade and there being a non-zero objective probability that I am a brain in a vat having a mere appearance of a table. This is significant since—as Hawthorne notes in the lately quoted passage—many philosophers are comfortable maintaining that we know that we are not brains in vats. Accordingly, these philosophers should feel comfortable maintaining that we know that certain particles have not undergone bizarre quantum rearrangements so as to compose a table façade. I do not claim that these philosophers are right to say that we know that we are not brains in vats. My point here is just that a sceptic who employs a case like the Quantum Case enjoys no dialectical advantage over a sceptic who adverts to more traditional sceptical hypotheses. Lottery-style puzzles—at least ones that should begin to trouble us such as the Quantum Case—do not require novel non-sceptical replies, or so it seems to me.

I have only scratched the surface of Hawthorne's impressive book. *Knowledge and Lotteries* is a philosophical page turner and it is chock full of many other novel ideas and arguments that I have not discussed here. The book is also a delightful read. Throughout, Hawthorne magnificently blends substance and rigour with stylistic flair. *Knowledge and Lotteries* will no doubt prove to be a significant contribution to epistemology in general and to the growing literature on the epistemology of lottery-style paradoxes in particular.