Hawthorne's Lottery Puzzle and the Nature of Belief*

Christopher Hill and Joshua Schechter

Brown University

Suppose that that T is one of ten million tickets that have been sold in a fair lottery. Clearly the probability that T will lose is extremely high. Suppose that I know these facts. Suppose, too, that T will in fact lose. Given these assumptions, it is tempting to conclude that I can know P, the proposition that T will lose. Certainly the present authors are tempted to conclude this. In his brilliant book *Knowledge and Lotteries*, however, John Hawthorne makes a strong case for the opposing view. He urges that in everyday contexts speakers are reluctant, or even outright unwilling, to claim that agents know propositions like P. And he buttresses this appeal to conversational data with several powerful theoretical arguments.

Hawthorne uses the expression "lottery proposition" to stand for propositions that satisfy two conditions: first, they have a very high degree of probability; and second, there is an intuitive reluctance to say that they are known to be true. He maintains that many propositions concerning lotteries meet these two conditions, and that a variety of other propositions meet them as well, including the proposition that Hawthorne's car has not been stolen since he parked it in a certain lot this morning, and the proposition that Hawthorne will not have a major heart attack in the near future.[1] It is clear that these propositions are highly probable. But also, according to Hawthorne, there is considerable intuitive reluctance to claim that they are within our ken. Moreover, Hawthorne maintains

---

* This is a slightly expanded version of a paper forthcoming in *Philosophical Issues.*

[1] The term "lottery proposition" and the claim that lottery propositions are widespread are due to Vogel (1990).

that this intuitive reluctance is well-founded. In most cases, he says, it is wrong to claim that an agent knows a lottery proposition.

We feel that Hawthorne's work on lottery propositions is of the first importance. He makes a case for his views about lottery propositions that is *prima facie* quite persuasive. Moreover, his views have important consequences. For example, as he shows, doubts concerning the knowability of lottery propositions provide motivation for a fairly general skepticism about our ability to know propositions of other kinds. There is also a further reason for valuing Hawthorne's discussion. He develops his views about lottery propositions in conjunction with some independently motivated ideas about the metaphysical relationships between knowledge and such things as explanation, assertion, epistemic possibility, epistemic probability, practical reasoning, theoretical reasoning, and practical interests.[2] These ideas have a strong appeal. Moreover, if they are sound, then traditional epistemology has failed to appreciate the metaphysical importance of knowledge. The ideas imply that knowledge plays a metaphysically foundational role with respect to communication, action, and reasoning.

Although we feel challenged and stimulated by Hawthorne's discussion, and feel grateful to him for bringing a wealth of new data and new forms of argumentation to the fore, we find ourselves in disagreement with many of his conclusions. We doubt that the intuitive reluctance to attribute knowledge of lottery propositions is as strong as he supposes. We also think that it is possible to explain much of the reluctance that actually exists in a way that neutralizes its probative value. Further, we have reservations concerning his claims about the metaphysical relations linking knowledge to other things. As we see it, these relations are both less systematic and less intimate than Hawthorne

[2] Many of these connections are also defended in Williamson (2000).

maintains. In this paper, we make a case for these dissenting opinions. We also put forward several positive proposals, urging, among other things, that assertion is governed by a Gricean constraint that makes no reference to knowledge, and that practical reasoning has more to do with rational degrees of belief than with states of knowledge.[3]

I

In motivating the claim that thinkers ordinarily do not know lottery propositions, Hawthorne relies on a line of thought that he calls *parity reasoning*. (pp. 4–5)[4] We will begin by reviewing this part of his discussion.

Consider a fair lottery with 1,000 tickets. Suppose the ticket that will in fact win is ticket #1000. Suppose that I possess ticket #1, and that tickets #2 through #999 are salient to me for some reason, perhaps because each is owned by a friend of mine. Suppose that I form a belief that my ticket will lose on the grounds that it has a 999/1000 chance of losing. Suppose, too, that I have a lot of time on my hands, and that to fill the hours I form beliefs about the tickets owned by my friends. In particular, for each of the tickets #2 through #999, I form the belief that it will lose. My evidence for each of these beliefs is exactly the same as my evidence for the belief that my own ticket will lose. In each case, my belief is based on purely statistical grounds.

Given this scenario, suppose that I count as *knowing* that my ticket, ticket #1, will lose. Since the grounds for each of my 999 beliefs are exactly the same, if one of these beliefs counts as knowledge, presumably so do the rest. Thus, I know that ticket #2 will

[3] The view we present here is a version of what Hawthorne calls “simple moderate invariantism” in section 3.7 of Hawthorne (2004).

[4] Unless otherwise indicated, page numbers in the text refer to Hawthorne (2004).

lose, I know that ticket #3 will lose, ..., and I know that ticket #999 will lose. But if I know that *each* of these tickets will lose, then surely I am in a position to know that *all* of these tickets will lose. This is an application of the principle that knowledge can be extended by conjoining propositions that one knows. Moreover, surely I am also in a position to know that ticket #1000 will win. This is an application of the more general principle that knowledge can be extended by deduction.

However, as Hawthorne notes, these two consequences seem preposterous. We do not think it possible to know by statistical reasoning that the first 999 tickets will lose – after all, the chance that one of them will win is very high. We also do not think it possible to know by statistical reasoning that ticket #1000 will win – ticket #1000 is just as likely to lose as any of the rest. So something must have gone wrong in the reasoning above. The natural suggestion is that the problem lies with the supposition that I know that ticket #1 will lose. Thus, Hawthorne argues, there is pressure to think that we do not ordinarily know such lottery propositions.

The argument can be generalized. For any lottery proposition – for example, for the proposition that my car has not been stolen since I parked it this morning – an analogous line of thought can be used to motivate the claim that I do not know it. Indeed, Hawthorne suggests that similar reasoning is what explains the responses ordinary thinkers have to lottery situations. When thinkers find it intuitively plausible that they do not know lottery propositions – and disavow such knowledge – this is typically because they (perhaps implicitly) engage in reasoning similar to the line of thought above.

Hawthorne characterizes this general pattern of reasoning as follows:

> *Parity Reasoning*. One conceptualizes the proposition that *p* as the proposition that one particular member of a set of subcases (p_1, …, p_n) will (or does) not obtain, where one has no appreciably stronger reason for thinking that any given member will not obtain. Insofar as one reckons it absurd to suppose that one is able to know of each of (p_1, …, p_n) that it will not obtain, one then reckons oneself unable to know that *p*. (p. 16)

According to Hawthorne, there are two reasons that thinkers may find it absurd to suppose that they are able to know of each of (p_1, …, p_n) that it will not obtain: It may be obvious that at least one of (p_1, …, p_n) will obtain, or it may be obvious that some consequence of the individual claims (such as their conjunction) cannot be known.[5]

Parity reasoning thus plays an important role in Hawthorne's monograph. He engages in it to motivate the claim that thinkers frequently do not know lottery propositions. He also appeals to it to explain ordinary thinkers' intuitions and responses to lottery cases. If we can find a principled basis to reject this general line of thought, that will go a long way toward allowing us to comfortably accept that thinkers ordinarily know lottery propositions. And it will substantially decrease the pressure to modify our general picture of knowledge.

Before we present our response to Hawthorne's parity argument, however, we should first note that there is something both plausible and illuminating in Hawthorne's discussion. It is appealing to think that parity reasoning is what explains the reactions

[5] In what follows, we will concentrate on parity reasoning involving the second reason. This is Hawthorne's focus (e.g., in his initial parity argument and on p. 179). Our response to the first reason will be implicit in what follows – it is possible to justifiably believe of each of (p_1, …, p_n) that it will not obtain while still knowing that at least one of them will obtain, although it is very natural for thinkers to make a mistake about this fact.

thinkers sometimes have to lottery cases, when they find it intuitive that they do not know lottery propositions and disavow such knowledge.[6] However, this reliance on parity reasoning does not show that such reasoning is correct. Indeed, we believe that parity reasoning is fallacious.[7]

It will not be surprising that we reject Hawthorne's parity argument on the grounds that the principles that one's knowledge can always be extended by conjoining known propositions – and, more generally, by deduction – are false.

Hawthorne carefully states the general principle as follows:

> *Multi-Premise Closure* (MPC). Necessarily, if S knows p_1, …, p_n, competently deduces q, and thereby comes to believe q, while retaining knowledge of p_1, …, p_n throughout, then S knows q. (p. 33)

The narrower principle can be stated analogously:

> *Conjunction Introduction* (CI). Necessarily, if S knows each of p_1 and p_2, competently deduces their conjunction, and thereby comes to believe their conjunction, while retaining knowledge of each of p_1 and p_2 throughout, then S knows their conjunction.

[6] We suspect that there are additional reasons ordinary thinkers find it intuitive that they do not know lottery propositions. For instance, in reflecting about knowledge, ordinary thinkers typically focus on a highly restricted class of paradigm cases of knowledge – simple mathematical and logical truths, truths about the immediate environment, and so on. Insofar as the epistemic status of lottery propositions is dissimilar to that of the paradigm cases, ordinary thinkers will be leery of attributing knowledge of them.

[7] Interestingly, Hawthorne also suggests that parity reasoning is fallacious, since it is a mistake to think that if someone knows *p* (in a given practical environment) then that person is in a position to use *p* as a premise in every practical environment. (p. 179) We agree. But we do not think that this is because what one knows depends on one's practical environment. Rather, it is because there is only a loose connection between knowledge and practical reasoning.

The main difficulty with these principles is familiar: Conjunction – and deduction from multiple premises more generally – can aggregate risk. A thinker may be justified in believing each of a set of propositions to a high degree but not in believing their conjunction to nearly as high a degree. For instance, using the probability calculus as a model of graded belief – admittedly an imperfect one – if two propositions each have probability .9, their conjunction may have a probability as low as .8. If, as seems plausible, knowing *p* is compatible with *p* having an epistemic probability lower than 1, MPC and CI will be subject to counterexamples.

The general phenomenon of risk aggregation is well-known from discussions of the preface paradox.[8] One version of this problem goes as follows: There are very many propositions that I count as knowing. Such propositions include simple claims of mathematics and logic; claims about myself, my environment, and my past experiences; and so on. Consider the conjunction of all of these claims. It seems that even were I cognitively able to competently deduce the conjunction from each of the individual claims that I know, I would not be in a position to know the conjunction. For I know that I sometimes – very rarely, perhaps – make mistakes. It should seem likely to me that at least one of the relevant propositions is false. It would be the height of arrogance to go on and infer the conjunction, knowing full well that there is a significant risk of falsity. As it happens, since the conjunction is a conjunction of claims that I *know*, if I were to draw

[8] See Makinson (1965). See Christensen (2004, chapter 3) for a recent discussion. For Hawthorne's discussion see pp. 48–9, 182–3.

the inference, I would infer a truth. But this would be a happy accident. The belief would not count as genuine knowledge. CI and MPC therefore fail.[9]

Of course, to address the parity argument, it is insufficient to simply reject MPC and related principles. We must also explain their appeal. Why it that we find it intuitive that, as Hawthorne puts it, “one can add to what one knows by deduction from what one knows”? (p. 46)[10] Absent an answer to this challenge, the response to the parity considerations is incomplete.

It will not be possible for us to address this issue until after we have developed our general picture of knowledge, communication, and reasoning. Let us therefore put aside this question for now and consider some of the striking claims Hawthorne makes about the relationships between knowledge, assertion, epistemic possibility, and practical reasoning.

II

So far we have been concerned only with the initial stages of Hawthorne’s discussion of his central claim, the doctrine that agents cannot ordinarily be said to know that lottery propositions are true. As we have seen, he urges that there is direct, intuitive support for this doctrine. He also maintains that the supporting intuitions are explained and buttressed by a line of thought that he calls the parity argument. But this is not the whole story. In addition to invoking our intuitions and considerations of parity, he provides the materials for constructing three additional arguments. Hawthorne does not state these

[9] Notice that cases like this do not put any pressure on Single Premise Closure.
[10] See Williamson (2000, p. 117) for a similar intuitive formulation.

arguments explicitly. Rather, he supplies the key premises and leaves the rest to the imagination of the reader. We wish to engage with these arguments because we would like our assessment of Hawthorne's central claim to make contact with all of the relevant considerations

We will call the three additional arguments the *assertion argument*, the *epistemic possibility argument*, and the *practical reasoning argument.*

The assertion argument begins with the claim that we are generally unwilling to assert lottery propositions. "Despite having good reason to think that a lottery ticket will lose, it is typically out of place to declare outright 'He won't win the lottery' in advance of the drawing and without special insider information." (p. 21) Hawthorne also endorses the increasingly popular doctrine that knowledge is the norm of assertion.[11] That is, he maintains that it is plausible, quite apart from considerations having to do with lotteries, that conversation is governed by the following principle:

(P1) A speaker ought not to assert that *p* unless he or she knows that *p.*

In Hawthorne's words: "The practice of assertion is constituted by the rule/requirement that one must assert something only if one knows it. Thus if someone asserts *p*, it is proper to criticize that person if she does not know that *p*." (p. 23)

Given (P1), it is possible to construct a best explanation argument for Hawthorne's view about knowledge of lottery propositions that runs as follows: "Suppose that we don't normally know that lottery propositions are true, and that speakers are generally aware of this fact. In combination with (P1), this assumption enables us to explain the

[11] See Williamson (2000, chapter 11).

fact that speakers are not normally inclined to assert lottery propositions: they don't assert lottery propositions because they recognize that they don't know that the propositions are true. Thus, the assumption that we don't know lottery propositions enables an explanation of a considerable body of data. Moreover, it is the best explanation, since other explanations can be shown independently to be inadequate." (Hawthorne does in fact consider two other explanations critically. Both are based on Gricean maxims.)

We turn now to the epistemic possibility argument. Hawthorne shares the familiar view that there is a connection between epistemic possibility and knowledge that can be expressed as follows:

> (P2) It is possible that *p* for S at t iff *p* is consistent with what S knows at t. (p. 26)

Now Hawthorne thinks that there is a use of "There is a chance that *p* is true" on which it is equivalent to "It is epistemically possible that *p*." Because of this, he holds that (P2) has a companion that can be formulated as follows:

> (P3) There is a chance of *p* for S at t iff *p* is consistent with what S knows at t. (p. 26)

In defense of these principles, Hawthorne claims that their validity is the best explanation of the fact that sentences of the form “I know that *p* but it is possible that not-*p*” and “I know that *p* but there is a chance that not-p” are not typically assertible.

In view of (P3), it seems to be possible to reason as follows: “If S knows the lottery proposition that ticket #666 will lose, then, given (P3), it cannot be true that there is a chance for S that ticket #666 will win. But it is clear that there *is* a chance for S that ticket #666 will win. After all, we are assuming throughout that lottery beliefs are based on probabilistic reasoning. S knows the relevant probabilities, at least to an approximation, and therefore knows that there is a real though small chance that ticket #666 will win.”

The practical reasoning argument is the most interesting of the three additional lines of thought. Hawthorne opens by claiming that it is normally inappropriate to use lottery propositions as premises in practical reasoning. He illustrates this claim by the following argument, which he quite rightly finds unacceptable:

> 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’d better sell the ticket.

Intuitively, he says, this argument is absurd *because* it uses a lottery proposition as a premise.

Having made this initial claim, Hawthorne goes on to endorse the view that knowledge is the norm of practical reasoning. He spells this view out in terms of the following principles, which he takes to be plausible independently of considerations having to do with lotteries:

> (P4) If one knows that *p*, then it is acceptable to use the proposition that *p* as a premise in one's practical reasoning – i.e., in one's deliberations about how to act. (p. 30)

> (P5) If one doesn't know that *p*, then it is not acceptable to use the proposition that *p* as a premise in one's practical reasoning. (p. 30)

In combination with Hawthorne's initial claim, these principles lead to two independent but closely related arguments for the conclusion that we cannot normally be said to know that lottery propositions are true.

The first argument is quite simple. It proceeds directly from the initial claim and (P4) to the desired conclusion. It is deductively valid.

The second argument has the form of an inference to the best explanation. If we assume that we do not normally know lottery propositions to be true, we can explain the unacceptability of the forgoing argument about selling one's lottery ticket by appealing to (P5). That this explanation is better than alternatives is shown by the fact that it accords with our intuitions about what is wrong with the argument: "It is clear that if one asks ordinary folk why such reasoning is unacceptable, they will respond by pointing out that the first premise is not known to be true." (pp. 29–30)

III

We turn now to the task of assessing these arguments. We will try to show that they all have disabling flaws.

We begin by observing that the first premise of the assertion argument, the claim that a speaker must know that *p* in order to have a conversational entitlement to assert that *p*, appears to have counterexamples. If a speaker has good reason to believe that *p* at the time of the assertion, then the speaker is not usually subject to rational criticism even if it turns out that *p* is false. That is to say, as far as we can see, there is no practice of criticizing such speakers for having spoken inappropriately, or for having exceeded the bounds of conversational entitlement. The falsity of *p* prevents agents from knowing *p*, but it does not by itself undermine entitlements to assert *p*. Similarly, it seems that it is conversationally quite appropriate for someone who has good evidence for *p* to assert that *p*, even if, unbeknownst to him, there are Gettier factors that prevent his justified belief from counting as knowledge.[12]

In proposing that knowledge is the norm of assertion, Hawthorne is in effect endorsing (P1):

(P1) A speaker ought not to assert that *p* unless he or she knows that *p*.

The forgoing considerations suggest that (P1) is incorrect. It is not immediately clear what should be adopted in place of (P1). It is worth noting, however, that it is easy to find

[12] See Kvanvig (forthcoming) for a similar objection.

a principle that is at least immune to counterexamples of the forgoing sort. This is true, for example, of (P1*):

> (P1*) A speaker ought not to assert *p* unless he or she believes *p* and is justified in so believing.

We do not wish to suggest that (P1*) is the whole truth about assertion, or even that it is wholly true. We think any account of assertion that must incorporate a number of diverse components. We cite it only to underscore the point that it is easy to accommodate the counterexamples to (P1). There is no need to try to "save" (P1) by resorting to heroic measures.

Our objection to the assertion argument is an obvious one, and it seems likely that Hawthorne is aware of it. He is probably presupposing the response to the objection that is presented in Williamson's *Knowledge and its Limits*, a book that Hawthorne cites as an inspiration for his own treatment of assertion. (p. 21) Williamson acknowledges that we generally refrain from criticizing speakers who assert propositions that they do not know, provided that they are justified in believing that the propositions are true, but he tries to explain this fact in a way that is ultimately consistent with the knowledge account of assertion. In giving this explanation, he begins by observing that a speaker who is justified in believing a proposition is generally (though not always) justified in believing that he knows the proposition. In view of this fact, Williamson claims, it is generally appropriate to see speakers as *trying* to conform to the norm linking assertion to knowledge, and as therefore deserving to be treated with indulgence when they fail. That is to say, it generally makes sense for us to *excuse* speakers who fail to conform to the norm. This is the core of Williamson's explanation. But he tries to buttress it by pointing out that it is often our way to allow minor violations of norms to go unpenalized. We tend to reserve criticism for violations that are particularly salient or particularly serious. "[T]he knowledge account does not imply that asserting *p* without knowing *p* is a terrible crime. We are often quite relaxed about breaches of the rules of a game which we are playing. If the most flagrant and most serious breaches are penalized, the rest may do little harm."[13]

Williamson's defense of the knowledge account is ingenious, but we feel that it falls short of success. A practice that involves a rule linking assertion to knowledge but allows minor violations of the rule is more complex than a practice that is based on a less demanding rule, such as the foregoing rule linking assertion to justified belief. In a practice of the former sort there will have to be a convention allowing excuses together with one or more rules specifying the gravity of various forms of infraction. Moreover, each participant in such a practice will have to keep track of the various psychological and epistemic factors that determine whether particular infractions should be excused. Thus, insofar as there is a distinction to be drawn between being justified in believing a proposition and being justified in believing that one knows a proposition, participants will have to discriminate between these two states in assessing the performance of speakers. And they will also have to assess the gravity of infractions. In view of these considerations, the justified belief account of assertion enjoys a *prima facie* advantage over the knowledge account. It is significantly simpler.

[13] Williamson (2000, p. 258).

Why then are Williamson and Hawthorne so taken with the knowledge account? As far as we can determine, they are moved largely or even entirely by two considerations. First, they believe that the knowledge account provides the best explanation of a certain variant of Moore's Paradox – that is, of the fact that it would be odd or "paradoxical" for someone to assert a proposition of the form "*p* but I don't know that *p*."[14] Second, they think (i) that speakers are generally unwilling to assert lottery propositions, and (ii) that the knowledge account provides the best explanation of this fact.[15]

We have strong reservations about this rationale. The knowledge account is not the only explanation of the Moorean Paradox. There are other, competing explanations that are at least as plausible.[16] Moreover, while (i) receives a certain amount of support from

[14] Williamson (2000, p. 253).

[15] See Williamson (2000, p. 249) and Hawthorne (2004, pp. 21–23). The second motivation might seem problematic in the current context. The fact that we don't know lottery propositions is supposed to gain support from the knowledge account of assertion. How can it also be used to support the knowledge account? The answer is that Williamson and Hawthorne should be seen as providing a complex explanation for our (alleged) unwillingness to assert lottery propositions. The explanation is the conjunction of the knowledge account and the hypothesis that speakers are aware that they do not know lottery propositions.

[16] For instance, it is possible to explain the inappropriateness of asserting "*p* but I don't know that *p*" using principle (P1*), which appeals to justified belief rather than knowledge. Suppose it is true that one ought not to assert a proposition unless one believes the proposition and one is fully justified in believing it. Suppose also that a certain speaker, Max, believes and is fully justified in believing the first conjunct of "*p* but I don't know that *p*." Will Max also believe and be fully justified in believing the second conjunct? For this to be true, it would have to be true that Max believes and is justified in believing at least one of the following four propositions: (i) Max does not believe that *p*; (ii) it is not true that *p*; (iii) Max is not justified in believing that *p*; and (iv) Max fails to satisfy the Gettier condition with respect to *p*. (Here of course we are assuming the standard theory of knowledge.) Could it be true that Max believes (i) and is justified in so believing? This is extremely unlikely. By hypothesis, Max believes that *p*. Given this hypothesis, it is very unlikely that Max believes (i) and is justified in believing it. Our assumptions also make it unlikely that Max believes (ii) and is justified in believing it. We have assumed that Max believes *p* and is justified in believing *p*. Given this assumption, it is much more likely that Max believes and is justified in believing that *p* is true than that he believes and is justified in believing that *p* is false. Further, as reflection shows, our assumptions also make it extremely unlikely that Max believes and is justified in believing (iii). What about (iv)? Well, if Max believed that he failed to satisfy the Gettier condition, and was justified in so believing, then he would not be fully justified in believing that *p*, which by hypothesis he is. To see this, observe that if one believes that he is in a Gettier situation with respect to a certain proposition, then he believes that special factors are present which prevent his *prima facie* justification for believing the proposition from being a satisfactory guide to truth. The latter belief is a defeater for one's *prima facie* justification for believing the proposition, so it precludes one's having a full or all-things-considered justification for believing it.

the data involving assertion, there is a large range of data that it fails to accommodate. It is often entirely appropriate to assert a lottery proposition.

In particular, there are many contexts in which it is entirely appropriate for agents to assert that particular lottery tickets will lose. It is also appropriate in such contexts for agents to assert that they *know* that particular tickets will lose. Suppose, for example, that Dan is trying to figure out whether he can afford to buy a sailboat. Dan might say, "Well, I have a lottery ticket. I *might* win the lottery. So I'd better factor that possibility in somewhere." Here it seems entirely reasonable for Dan's friend Jerry to respond as follows: "Get serious. We both know that you're not going to win the lottery. You should just forget about that possibility." It is clear that conversations of this sort happen all the time.[17]

It appears, then, that the assertion argument is problematic. It is based on an account of assertion that is *prima facie* incompatible with a large array of data. To be sure, as Williamson has shown, there is a way of qualifying the account so as to accommodate this data; but Williamson's proposal incurs substantial theoretical costs. Further, the knowledge account of assertion has very little positive motivation. As far as we can determine, its primary support comes from a version of Moore's Paradox and the perception that speakers are reluctant to assert lottery propositions. But there are other accounts of assertion that can claim support from Moore's Paradox, and the perception

---

To summarize: The justified belief account of assertion implies that Max must satisfy a belief condition and a justification condition with respect to *p* in order to be entitled to assert "*p* but I don't know that *p*." It also implies that he must satisfy a belief condition and a justification condition with respect to "I don't know that *p*" in order to be entitled to assert this proposition. It appears, however, that if Max satisfies these conditions with respect to the first conjunct of "*p* but I don't know that *p*," he will almost certainly fail to satisfy them with respect to the second conjunct. This explains the inappropriateness of asserting the proposition.

[17] Indeed, Hawthorne presents a case of this sort. (p. 84)

concerning lottery propositions is only partially correct. There are plenty of occasions on which speakers find it entirely natural and appropriate to assert lottery propositions.

By themselves, these considerations do not amount to a knockdown objection to the knowledge account, nor to a decisive reason for rejecting the assertion argument. Perhaps it is possible to accommodate the complexities in the data involving lottery propositions by introducing additional qualifications. It is clear, however, that the foregoing considerations provide reasons for being dissatisfied with the knowledge account, and for hoping that it is possible to provide a better explanation of the data involving assertion. In fact, we think that it is possible to do better. With a view to make this plausible, we will propose a new account of assertion – an account that explains the comparatively simple conversational data involving ordinary propositions, and that also does a better job than the knowledge account of explaining the complex conversational data involving lottery propositions. We begin the process of formulating this account in the next section.

IV

If agents are often willing to assert lottery propositions, why does it seem otherwise to Williamson and Hawthorne? Perhaps it seems otherwise because there is a type of situation in which it is clearly inappropriate to give the advice that Jerry gives in the foregoing sailboat example. Situations of this sort are by no means the norm; but they are familiar to everyone, and they have a certain salience. We will describe these situations in the present section, and will also offer the beginnings of an argument that they are best explained in terms of a Gricean principle governing the flow of information. Our Gricean explanation is incompatible with the explanation favored by Williamson and Hawthorne, which is of course based on the knowledge account of assertion. Accordingly, we see our Gricean explanation as driving the final nail into the coffin of the knowledge account.

The situations we have in mind arise when an agent A_1 is involved conversationally with an agent A_2 who is considering a course of action that might lead to a considerable gain or a considerable loss, depending on whether a certain lottery proposition is true. Suppose, for example, that Carol is wondering whether to assert that Mary's lottery ticket will lose, and suppose also that Mary is trying to decide whether to hold onto the ticket or to sell it to a third party for a penny. In this case, Carol would very likely refrain from asserting the proposition, even though she has, and recognizes that she has, an excellent reason for believing that the proposition is true. The reason for this reluctance, we submit, is the Gricean principle that when one asserts that *p*, one thereby implicates one's belief that the proposition that *p* is more relevant to the audience's informational needs than any of the other propositions that one is justified in believing.[18] Thus, if Carol were to assert the proposition that Mary's ticket will lose, she would be implicating that she believes that this proposition is more relevant to Mary's present informational needs than any other proposition about Mary's ticket that she believes, such as the proposition that there is a very small but nonetheless positive probability that Mary's ticket will win. In fact, however, as we will explain in detail later on, the latter proposition is much more relevant to Mary's present informational needs than the former proposition, because it is the latter proposition that she should be using as a premise in deliberating about whether to hang onto her ticket or sell it for a penny. It can be assumed that Carol is aware of the

[18] This principle is not one that Grice put forward, though it fits well with his general view of conversation as a cooperative enterprise. It also has affinities with his maxim of relevance. See Grice (1975).

nature of Mary's current deliberations. By the same token, it can be assumed that she is aware that the proposition that there is a very small probability that her ticket will win is more relevant to Mary's present informational needs than the proposition that her ticket will lose. Since she is aware of these things, she must also be aware that it would be misleading for her to say something that would implicate that her most relevant belief is the belief that Mary's ticket will lose.

In section VII we will attempt to explain and justify our claim that it is the proposition that there is a very small probability that Mary's ticket will win, and not the proposition that the ticket will lose, that Mary should use as a premise in her practical deliberations, and also the attendant claim it is the former proposition that is most relevant to Mary's present informational needs. It is enough for our present purposes that the reader agree that *if* these claims are correct, then it would be inappropriate for Carol to assert the proposition that Mary's ticket will lose. We hope that we have made this conditional claim plausible.

V

We turn now to the epistemic possibility argument. It can be formulated as follows:

> *First premise*: For all *p*, S, and t, there is a chance of *p* for S at t iff *p* is consistent with what S knows at t.

> *Second premise*: Where A is an agent, L is the proposition that A's lottery ticket will win, and T is a time prior to the drawing, there is a chance of L for A at T.
>
> *Lemma*: Hence, L is consistent with what A knows at T.
>
> *Third premise*: If L is consistent with what A knows at T, then A does not know not-L at T.
>
> *Conclusion*: A does not know not-L at T. That is, A doesn't know that A's ticket will lose.

The second premise holds in virtue of A's being aware that there is one chance in ten million that the ticket will win.

We claim that the argument is fallacious. To be specific, there is an equivocation on the word "chance." In our view, the first premise is a conceptual truth, a principle that holds because it is *constitutive* of the concept of there being a chance for an agent that a proposition is true. In other words, the relationship between the notion of epistemic possibility and the notion of knowledge is of the same sort as the relationship between the notion of logical possibility and the notion of logical necessity: they are interdefinable. Accordingly, the relevant notion of chance has no content over and above the notion of being logically consistent with a certain body of knowledge. It has no special connection either with the concept of the objective probability of an event or the subjective probability of a proposition. On the other hand, "chance" has a quite different meaning in the second premise. As we observed a moment ago, the second premise holds in virtue of A's being aware of the objective probability that his ticket will win. The notion of chance here is an essentially probabilistic notion, a notion that has no special relationship with the notion of logical consistency or the notion of an agent's body of knowledge. It follows that the argument involves a fallacy of equivocation. The first and second premises lack the mutual relevance that is required for successful inference.

Our rationale for viewing the argument in this way has two parts. First, in explaining the notion of chance that figures in the first premise, Hawthorne tells us that it comes to the same thing as the notion of epistemic possibility. Now as many authors have pointed out, it is extremely plausible that the notions of epistemic possibility and knowledge are interdefinable. In particular, whether a proposition is epistemically possible for an agent is a matter of whether the proposition is logically consistent with the agent's body of knowledge. In view of this, there is good reason to think that this concept has no essential relation either to the concept of objective probability or to the concept of subjective probability. Second, if the second premise of the epistemic possibility argument made use of the same concept of chance as the first premise, it would be necessary to know the composition of A's body of knowledge in order to determine whether the premise is true. More specifically, it would be necessary to know whether the negation of L is a component of A's body of knowledge. But this means, in effect, that it would be necessary to know whether the negation of L is known by A. Any claim to possess such knowledge would beg the question, for the point of the argument is to establish that A does not know that L is false.

Thus far our evaluation of the epistemic possibility argument has presupposed a rather simple picture of the notion of chance that figures in the first premise. According to this simple picture, the relevant notion of chance is definable in terms of the concept of knowledge and the concept of logical consistency. Reflection shows that this picture fails to capture the full complexity of the notion. The notion of chance that is definable in terms of knowledge is best seen as involving probability rather than logical consistency. That is to say, when we speak of a proposition as having a certain probability in relation to knowledge, what we mean is that the proposition has a certain conditional probability relative to the class of propositions that are known to be true.[19] Now of course, when we are speaking of this form of probability, it is true to say that propositions that are known to be true have the highest degree of probability, and by the same token, it is true to say that the negations of these propositions have a probability of 0. Accordingly, when we have this form of probability in mind, it is inappropriate to assert a sentence of the form "S knows that *p* but there is a (non-zero) chance for S that not-*p*." But this should not blind us to the fact that it would be natural and appropriate to assert a sentence of the same form in circumstances in which it were clear that some other form of probability

---

[19] It is natural to use the expression "epistemic probability" to stand for this form of probability; but this usage can lead to serious confusions, for it is also natural to use "epistemic probability" as a term for the degree of belief that it is *rational* for an agent to assign to a proposition. To repeat, the first form can be explained as the conditional probability of a proposition relative to the class of propositions that an agent knows to be true. The second form can be explained as the probability that a proposition possesses relative the *evidence* that an agent possesses. Unless one shares Williamson's view that an agent's evidence is identical with the class of propositions that the agent knows to be true, one will want to distinguish sharply between these two forms of probability. The term "epistemic" makes it hard to keep this distinction in view.

There are still other ways of using "epistemic probability." For instance, it can be used as a term for the degree to which an agent is *epistemically justified* in believing that *p*. This may or may not come to the same thing as using it to refer to rational degrees of belief.

In this paper we use the expression "epistemic probability" in a fourth way – as a term for whatever it is that we reason with when we reason with probabilities or degrees of belief.

was under discussion. For instance, it would be entirely appropriate to claim that S knows that *p* is true while acknowledging both that it is rational for S to assign a non-zero degree of belief to not-*p* and that not-*p* has a non-zero objective probability.[20]

To summarize: The epistemic possibility argument fails because it does not distinguish among the various forms of probability that are relevant to questions of knowledge.

VI

The practical reasoning argument actually consists of two lines of thought. The first is a deductive argument with two premises. One premise is the claim that it is generally inappropriate to use lottery propositions as premises in practical reasoning. The other is (P4):

> (P4) If one knows that *p*, then it is acceptable to use the proposition that *p* as a premise in one's practical reasoning – i.e., in one's deliberations about how to act.

In combination, these premises entail that we generally lack knowledge of lottery propositions. The second line of thought is a best explanation argument. Its first premise is identical with that of the first line of thought. The second premise is (P5):

[20] Indeed, it seems that it would be arrogant to assign the highest possible degree of belief to many of the propositions that we know.

> (P5) If one doesn't know that *p*, then it is not acceptable to use the proposition that *p* as a premise in one's practical reasoning. (p. 30)

Once these two premises have been stated, we are invited to observe that it is possible to explain why it is generally inappropriate lottery propositions in practical reasoning by invoking (P5) and the hypothesis that it is generally impossible to know lottery propositions. We are also told that this explanation is superior to all competing explanations.

There are problems with the premises of both arguments.

We have already considered a case, the sailboat example, in which an agent clearly has a right to use a lottery proposition in his practical reasoning. It would be easy to construct additional cases. Accordingly, the claim that serves as the first premise for both arguments is false. The claim enjoys a certain initial plausibility, so presumably some fraction of its content is correct. That is to say, there is presumably a circumscribed class of cases in which it is inappropriate to rely on lottery propositions in practical reasoning. But there is a significant range of cases in which just the opposite is true.

Moreover, it is clear that (P4) and (P5) are flawed. It is entirely possible for one to feel extremely confident that one knows a proposition while recognizing that it would be inappropriate to use the proposition as a premise in certain forms of practical reasoning. Suppose that Jill has excellent reason for holding that Jack is not in New York today – perhaps she saw someone who looks *exactly* like Jack on Benefit Street in Providence a short while ago, and the person in question responded to her wave with a broad smile of recognition. In these circumstances, Jill will feel that she has every right to claim that she knows that Jack is not in New York. But will she be willing to stake the lives of her children on this proposition? Of course not! Despite feeling sure that she knows that Jack is not in New York, she will recognize that there is a miniscule chance that she is wrong. (Here we remind the reader of the preceding section, in which we maintain that knowledge of *p* does not preclude one's allowing that there is a chance that not-p.) Because of this, Jill will think it inappropriate to accept certain bets involving the proposition – specifically, bets in which she stands to lose a great deal. Moreover, the rest of us will applaud this attitude. So (P4) is wrong. To see that (P5) is also wrong, observe that in certain cases, anyway, one is entitled to use a proposition in one's practical reasoning simply because one is epistemically justified in believing it. The proposition need not be true, and one need not satisfy whatever proposition it is that normally serves to block Gettier counterexamples. For example, if weather.com has promised beautiful weather, and Jane relies on this announcement in formulating her plan for the day, she will not be open to rational criticism if the announcement turns out to be false. People may say that her decision to spend the day at the beach was unfortunate, or even inadvisable, but they will not say that she went wrong in her practical reasoning.

We may conclude, then, that the practical reasoning argument does not work.

VII

But this should not be the end of our inquiry. It is clear that the practical reasoning argument fails, but we should consider whether there is an argument that fares better.

After all, Hawthorne is quite right in claiming that the following argument is unacceptable:

> 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'd better sell the ticket.

This shows that there is a type of situation in which it is inappropriate to use a lottery proposition as a premise. What are the common features of situations of this type? What distinguishes them from other situations? How frequently do they occur? Also, do such situations have any implications concerning knowledge of lottery propositions? Is it perhaps true that it is inappropriate to rely on lottery propositions in such situations because there is something about them that precludes knowledge of lottery propositions?

In considering these questions, it is important to keep in mind the distinction between probabilistic reasoning that is *certainty-based* and practical reasoning that is *probability-based* – or in other words, the distinction between *decision-making under certainty* and *decision-making under risk*. In a situation in which an agent is attempting to decide among a set of actions $A_1, \ldots, A_n$, it may be true that the agent is in a position to predict the outcome of each action with certainty, or it may be that for each A_i, the agent is forced to recognize two or more possible outcomes for A_i, and can do no more than assign a probability to each of them. It is relatively simple to choose an action in situations of the first sort: one need only perform a cost/benefit analysis, adding up the benefits that would accrue from each A_i and subtracting the costs, and then choose the action that would be the most beneficial. On the other hand, practical reasoning that is probability-based involves much more complicated calculations. In a case that calls for probability-based reasoning, it is necessary to take information about probabilities into account. It is not known exactly what form these probabilistic calculations take in actual human decision making, though there are some well motivated proposals, such as prospect theory. In the interests of definiteness, we will suppose that probability-based practical reasoning involves calculations of expected utilities and that choices are aimed at maximizing expected utility. This assumption is no doubt unrealistic in a variety of respects, but it has the value of illustrating the point that a procedure that is sensitive to a range of probabilities will be much more complex than a procedure that in effect treats all probabilities as 0 or 1.

Since practical reasoning that is certainty-based is much simpler than practical reasoning that is probability based, it behooves one to rely on the former whenever it is possible to do so. Indeed, the savings in time, energy, and memory is so great that it behooves one to act *as if* all probabilities were 0 or 1 in cases in which the actual probabilities are close to those extreme quantities, provided that doing so would not significantly change the final results. We will say that a situation in which it is appropriate to act as if low probabilities were 0 and high probabilities were 1 is a *simplification-permitting* situation. Thus, in a simplification-permitting situation, there is a simplifying idealization of probabilities that (a) enables one to use a certainty-based procedure for choosing an action rather than a probability-based procedure, and (b) yields the same conclusion concerning which action to perform as would a calculation that took information about actual probabilities into account.[21]

We now wish to claim that when probabilities are close to 0 and 1, it is generally true that one is in a simplification-permitting situation. One will be in a situation that is not simplification-permitting only if the costs or benefits associated with one or more of the outcomes are quite large.[22] In a case of this sort, the difference between a very low probability and 0, or the difference between a very large probability and 1, will generally be such as to have an impact on calculations of expected utility. But in other cases, the differences between realistic assumptions about probabilities and idealizations will not be reflected in the end result. Cases of this second sort will be the norm, for it is rare that an agent stands to gain or lose a great deal by the choice of a single action.

Of course, for the distinction between simplification-permitting and non-simplification-permitting situations to be of practical importance, it would have to be possible to identify simplification-permitting situations as such in advance of performing the calculations that probability-based reasoning requires. Fortunately, it generally *is* possible to recognize simplification-permitting situations in advance. This is illustrated by the forgoing sailboat case, and also by the Carol/Mary case. In the sailboat case it is reasonably clear in advance that the probability of Dan's winning the lottery is so low as

---

[21] In addition to situations in which it is appropriate to act as if *all* of the extreme probabilities were 0 or 1, there are also situations in which it is appropriate to act as if one or more of the extreme probabilities were 0 or 1, but to preserve the exact values of one or more of the other extreme probabilities. This happens when one or more, but not all, of the extreme probabilities are associated with states of affairs that have little intrinsic or instrumental value. A situation of this sort cannot be said to be simplification-permitting in the strong sense of the term that is defined above, but they are simplification-permitting in a weaker sense. In the interests of simplicity, we will not distinguish between these two types of simplification-permitting situation in the sequel. Much of what we will say applies to situations of both types. When this is not true, the context will make it obvious which type we have in mind.

[22] We assume here that the relevant set of possible outcomes for each action is relatively small. If there are very many possible outcomes, many with a very small probability of occurring, the situation will not be simplification-permitting. We leave this qualification implicit in what follows.

to be of virtually no consequence with respect the overall probability of Dan's having enough money to purchase a sailboat. That is why it is entirely natural for Jerry to urge Dan to forget about his ownership of a lottery ticket in considering his financial prospects. In the Carol/Mary case, on the other hand, both Carol and Mary should be able to see, in advance of grinding through the actual calculations, that the expected utility of hanging onto the ticket might not exceed the expected utility of selling the ticket, given that the latter quantity is just one cent. That is to say, in this case it is pretty clear in advance that Mary might come to a different conclusion concerning the disposition of the ticket if she were to apply probability-based reasoning than if she were to adopt the simplifying assumption that there is no chance of her winning.

To amplify: In the sailboat case Dan is considering whether to take his ownership of a ticket into account in determining whether he will be able to afford to buy a sailboat. It would clearly be expensive in terms of time and energy if he were to take his ownership of his ticket into account in his calculations. Moreover, since the probability of his winning is very low, it is reasonably clear in advance that the ultimate conclusion of his reasoning would not be materially affected by taking his ownership of the ticket into account. In view of these considerations, the costs associated with taking his ownership of the ticket into account in his calculations exceeds any gain that might accrue from taking it into account. By the same token, it makes sense for Dan to proceed on the assumption that his ticket will lose. In the Carol/Mary case, on the other hand, it is definitely *not* clear in advance that taking Mary's ownership of the ticket into account would have no effect on the calculations. Mary is considering the question of whether to hang onto her ticket or to sell it for a penny. Even a cursory look at the relevant factors, the probability of her winning and the amount of the prize, shows that the expected utility of her owning the ticket is highly pertinent to this question. Or at least, this will be true if, as is usual, the lottery is one in which the winner will receive a large amount of money. Suppose, for example, that Mary's ticket is one of 10,000,000, and that the winner will receive $5,000,000. Here one need not calculate the actual expected utility of her hanging on to the ticket to appreciate that it is probably greater than a penny, which is the expected utility of selling the ticket. (The expected utility of hanging onto the ticket is actually fifty times as large as the expected utility of selling it.) By the same token, it would be inappropriate for Mary simply to assume that her ticket will lose. That would be equivalent to assuming that the expected utility of continuing to own the ticket is zero.

We have thus far been concerned to develop a general theory of when and why it is appropriate to use lottery propositions as premises in practical reasoning, without taking their probabilities into account. We now observe that the theory provides answers to two of the questions abut lottery propositions that arose at the beginning of the present section. In the first place, it enables us to explain what it is that distinguishes situations in which it is appropriate to use a lottery proposition as a premise in practical reasoning, without taking its probability into account, and situations in which it is not. The former situations are precisely the ones that can be seen in advance to be simplification-permitting. To spell this out a bit, they are the situations in which it can be seen in advance that the actual probabilities are sufficiently close to 0 and 1, and the costs and benefits are sufficiently modest, that certainty-based practical reasoning will produce the same answers as probability-based reasoning. The theory also provides an answer the question of how frequently situations of the sort in question arise. The answer is "quite frequently." This holds for two reasons. First, as we have already seen, cases in which it is safe to rely on simplifying assumptions about the probabilities of lottery propositions are the norm. And second, it seems to be generally possible to determine in advance whether a situation is one in which it is safe to rely on simplifying assumptions. All one has to do, to make such a determination, is check whether the probabilities are close to 0 and 1, and the costs and benefits are modest.

In section IV we asserted the following conditional: *If* it is true that Mary should use the proposition that there is a very small probability that her ticket will win in her practical reasoning, and not the proposition that the ticket will lose, *and* it is therefore true that it is the former proposition that is most relevant to Mary's present informational needs, *then* it would be inappropriate for Carol to assert the proposition that Mary's ticket will lose. We hope that we have now succeeded in making it plausible that the antecedent of this conditional is true. If so, we can claim to have offered a reasonably complete diagnosis of what is wrong with Hawthorne's assertion argument. Contrary to what the argument contends, it is possible to explain a speaker's unwillingness to assert a lottery proposition without assuming that the speaker fails to know the proposition. When an unwillingness of this sort exists, it is possible to explain it by citing the Gricean duty not to assert propositions that are irrelevant to the perceived informational needs of the audience.

Perhaps it will be useful to summarize our explanation of why it is inappropriate for Carol to assert that Mary's ticket will lose. If Carol were to assert the proposition, she would be implicating her belief that the proposition is more relevant to Mary's current informational needs than any of the other propositions about Mary's ticket that Carol is justified in believing. Because of this, she would be misleading Mary if she were to assert the proposition, for it is in fact true, and apparent to Carol, that the proposition that there is a very slight chance that Mary's ticket will win is much more relevant to Mary's informational needs than the proposition that the ticket will lose.

VIII

The assertion argument, the epistemic possibility argument, and the practical reasoning argument are all ultimately unpersuasive. Let us now return to our earlier discussion of parity reasoning. Recall that such reasoning relies upon Multi-Premise Closure and related principles. We earlier rejected MPC on the grounds that deductive inference from multiple premises can aggregate risks. But we also noted that an adequate response to the parity argument requires providing an explanation of why we find MPC so intuitively compelling. A fully adequate response must also show why Hawthorne's response to aggregation of risk considerations fails, and why the main motivation he puts forward for MPC is problematic. We are now in a position to answer these challenges.

Let us begin by considering Hawthorne's response to the aggregation of risk considerations. Recall that we argued that MPC fails because deduction can aggregate risk in a way that destroys knowledge. In responding to this objection, Hawthorne relies on his discussion of chance. He suggests that knowing *p* is incompatible with there being any chance that *p* is false. Since risk is just the chance of falsity, in cases where thinkers know the premises of a deductive argument, there is no risk to aggregate. (p. 48)[23]

[23] Also see Williamson (2000, pp. 123–30).

As we have already seen, however, Hawthorne's discussion of chance trades on an equivocation on the word "chance". Knowing *p* may be incompatible with there being a chance that *p* is false, for the sense of "chance" definable in terms of knowledge and logical entailment (or conditional probability). But it is not incompatible with there being a chance *p* is false, for the sense of "chance" relevant to epistemic probabilities. And it is this latter notion that is relevant to risk aggregation. Hawthorne's response therefore fails.

Let us now turn to Hawthorne's primary positive motivation for MPC. He suggests that anyone who rejected MPC for a particular argument could be made to sound very much like Lewis Carroll's foolish Tortoise.[24] Such a person would presumably answer yes when asked whether he accepted each of the premises of the argument. He would also answer yes when asked whether he accepted that the premises jointly entail the conclusion of the argument. But he would answer no when asked whether he accepted the conclusion of the argument. And this seems to be an "exceedingly strange" pattern of responses. Insofar as we should avoid any commitment to such a pattern of responses, we should accept MPC. (pp. 39, 49)

This argument relies on the claim that knowledge is the norm of assertion. As we have already seen, that is an implausible view. What is more plausible is that in typical conversational contexts, a speaker will (and should) assert a claim only if he or she believes that the claim is relevant to the audience's informational needs. This Gricean principle can be used to explain why the above pattern of responses seems so strange, even assuming the falsity of MPC. Given the Gricean principle, it is natural to expect that speakers will only assert a sequence of premises followed by the claim that the premises jointly entail some conclusion if the speaker accepts the conclusion and takes the

[24] See Carroll (1895).

inference to be one that the audience should draw. The “exceedingly strange” pattern of responses violates this expectation. This explains why it would feel uncomfortable to answer the sequence of questions as above, and why it would sound so odd to hear such a pattern of responses. This discomfort would obtain notwithstanding the fact that the answers to the questions would all be true.

Finally, let us turn to the question of why MPC is such a tempting principle. Why do we find it intuitively compelling that we can add to our knowledge by drawing competent deductions from what we know? The general answer we want to propose is that in the most salient cases of deduction from known premises, MPC obtains. Our mistake is one of overgeneralization. Let us explain.

In our thinking about knowledge, we find it very natural to focus on certain paradigm cases. For example, we think about basic logical and mathematical knowledge, knowledge of simple conceptual truths, perceptual knowledge about the immediate environment, knowledge based on memory of the recent past, knowledge of well-known historical facts, and the like. These are the clearest and most central cases of knowledge. Ordinary thinkers are most comfortable attributing knowledge in such cases. Indeed, it is quite plausible that our understanding of the concept of knowledge is tied to our attribution of knowledge in these paradigm cases.

In each of these paradigm cases of knowledge, we are extremely confident about the truth of the relevant propositions. What could be more certain than our knowledge of simple logical and mathematical truths, and our perceptual knowledge about the immediate environment? In our reflections about knowledge, then, we focus on examples of known propositions with extremely high epistemic probabilities.

In our thinking about deduction, we find it natural to focus on deductive arguments similar to the ones we actually employ. Such arguments are relatively simple, with a small number of premises and inferential steps. This is because our minds are subject to significant cognitive limitations, including limitations on our short-term and working memory capacities. Indeed, we typically think only about particularly simple examples of deductive arguments. This is because we typically consider arguments that we can survey entirely in our heads. In our reflections about deduction, then, we focus on deductive inferences with only a few premises and inferential steps.

Taken together, these considerations suggest the following moral: In our thinking about knowledge and deduction, the cases of deductive argument from known premises that are salient to us are ones in which (i) the epistemic probability of each premise is extremely high and (ii) there are only a few premises and inferential steps.

There is a general fact about probabilities that is relevant here. Given any deductively valid argument, the probability of the conclusion will be at least as high as the probability of the conjunction of the premises. Therefore, if the probability of each premise is extremely high and there are only a few premises, the probability of the conclusion will be very high, too. For example, if there are two premises, each with probability .98, the probability of the conclusion will be at least .96. (It will be at least .9604 if the two premises are probabilistically independent.)

Using this general fact about probabilities, we can conclude that in cases of deductive arguments from known premises that are salient to us, the epistemic probability of the conclusion will be very high. It is plausible that knowing *p* entails justifiably assigning a high epistemic probability to *p*, but is compatible with not assigning 1 to *p*. Given this connection between knowledge and epistemic probability, in all – or almost all – of the salient cases of deduction from known premises, the conclusion of the deduction will count as knowledge. This is what explains the intuitive appeal of MPC. Since MPC holds in the salient cases of deduction from known premises, we find it natural to believe that it holds generally. And that is a mistake.[25]

Let us summarize our overall response to Hawthorne's discussion of parity reasoning. Parity reasoning is fallacious because it relies upon MPC or related principles. Such principles are false since deductive inference can aggregate risk in a way that can destroy knowledge. The intuitive appeal of MPC traces back to the fact that in *salient* examples of deduction from known premises, MPC obtains. The seductive mistake we make is in confusing this fact with the much stronger claim that in *all* examples of deduction from known premises, MPC obtains.

IX

Before we conclude, we would like to briefly point out a few of the major idealizations present in the picture of belief and reasoning we have been sketching.

We have already noted one idealization in our picture: it is extremely unrealistic to think that our practical deliberations closely resemble the explicit calculations of expected utilities familiar from decision theory. As we also noted, there are accounts in

[25] There are two other potential sources of the intuitive appeal of MPC worth mentioning. First, its appeal may be due in part to the fact that we frequently use mathematics as a model for knowledge. In mathematics, known logical entailment can generally be used to extend knowledge. Second, as we earlier argued, when a thinker assigns a high epistemic probability to *p* he or she can typically treat *p* as though it were certain, at least for the purposes of practical reasoning. We suspect that there is a similar phenomenon for theoretical reasoning, too: Thinkers can reasonably treat propositions with high epistemic probabilities as though they were certain for the purposes of theoretical reasoning. The correct way to characterize this phenomenon is somewhat murky, however. We hope to discuss surrounding issues on a future occasion.

the literature – such as prospect theory – that may do a better job of capturing human decision making.

There is, however, a still more fundamental issue. Throughout this paper, we have relied upon a notion of epistemic probability and have illustrated many of our points using the probability calculus. It is worth pointing out two important ways in which this is a simplification of our actual views.

First, we have been treating epistemic probabilities as though they were beliefs whose contents involve the attribution of (something like) probabilities to propositions. Correlatively, we have been treating decision-making under risk as though it always involved reasoning with such contents. We have adopted this familiar way of talking in order to simplify our exposition. But we would like to note that it is not essential to our general picture. It is possible to understand what we have been calling "epistemic probabilities" in terms of degrees of conviction or confidence. [26] On this view, decision making under risk is reasoning that takes into account differences in the degrees to which various propositions are believed. We stress this point because we have reservations about the view that decision making under risk typically involves reasoning with contents that attribute (anything like) probabilities. Perhaps in our explicit reasoning about lotteries, we reason about the objective probability of winning. But this is a rare phenomenon. Having differing degrees of confidence, in contrast, seems to be commonplace. [27]

Second, we have been treating epistemic probabilities – or degrees of belief – as though they obeyed (or ought to obey) the probability calculus. Here, too, we have severe

[26] See, for instance, Christensen (2004, chapter 2).

[27] We also find some plausibility in the idea that there are, in some sense, two distinct systems of belief – one for the kind of belief that is all-or-nothing and one for the kind of belief that comes in degrees.

reservations. We do not believe logical equivalents to the very same degrees. Nor would it be rational to always so believe. For instance, it would not be rational to have the highest degree of belief in a complex logical truth, at least not in the absence of a proof. A more refined account of rational degrees of belief seems required. Our picture of belief and reasoning, then, should be seen as committed to realism about degrees of belief. But it is not committed to the claim that degrees of belief must satisfy the probability calculus to count as rational.

X

Hawthorne's discussion of lottery propositions is rich and very stimulating. The arguments he puts forth for the claim that thinkers do not ordinarily know lottery propositions are arresting. However, as we have seen, they are ultimately unpersuasive. The direct, intuitive evidence is not nearly as compelling as Hawthorne supposes. The parity argument relies on a fallacious principle connecting knowledge and deduction. The assertion, epistemic possibility, and practical reasoning arguments rely on incorrect claims about the metaphysical connections between knowledge on the one hand, and assertion, chance, and practical reasoning on the other. As we have seen, once we adopt an appropriately complex view of the nature of belief and reasoning – and of the metaphysical connections between knowledge, communication, action, and reasoning – it becomes evident that Hawthorne has not yet made a case for his view of lottery propositions.

References

Carroll, Lewis. (1895) "What the Tortoise Said to Achilles." *Mind* 4: 278–280.

Christensen, David. (2004) *Putting Logic in its Place*. Oxford: Oxford University Press.

Grice, H. P. (1975) "Logic and Conversation." In *The Logic of Grammar*, edited by Donald Davidson and Gilbert Harman. Encino, CA: Dickenson, 64–75.

Hawthorne, John. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.

Kvanvig, Jonathan. (forthcoming) "Assertion, Knowledge, and Lotteries." In *Williamson on Knowledge*, edited by Duncan Pritchard and Patrick Greenough. Oxford: Oxford University Press.

Makinson, David. (1965) "The Paradox of the Preface." *Analysis* 25: 205–207.

Vogel, Jonathan. (1990) "Are there Counterexamples to the Closure Principle?" In *Doubting: Contemporary Perspectives on Skepticism*, edited by M. Roth and G. Ross. Dordrecht: Kluwer Academic Publishers, 13–27.

Williamson, Timothy. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.