Lotteries, Knowledge, and Irrelevant Alternatives

RACHEL MCKINNON *University of Calgary*

ABSTRACT: The lottery paradox plays an important role in arguments for various norms of assertion. Why is it that, prior to information on the results of a draw, assertions such as, "My ticket lost," seem inappropriate? This paper is composed of two projects. First, I articulate a number of problems arising from Timothy Williamson's analysis of the lottery paradox. Second, I propose a relevant alternatives theory, which I call the Non-Destabilizing Alternatives Theory (NDAT), that better explains the pathology of asserting lottery propositions, while permitting assertions of what I call fallible propositions such as, "My car is in the driveway."

RÉSUMÉ : Le paradoxe de la loterie joue un rôle important dans l'argumentation visant à défendre diverses normes de l'assertion. Comment se fait-il que, avant que les résultats d'un tirage soient connus, des assertions comme «Mon billet a perdu» semblent inappropriées? Cet article se compose de deux projets. Premièrement, je relève certains problèmes issus de l'analyse du paradoxe de la loterie par Timothy Williamson. Deuxièmement, je propose une théorie des alternatives pertinentes que j'appelle la «théorie des alternatives non-déstabilisantes» (NDAT), et qui explique d'une meilleure façon la pathologie de l'assertion de propositions concernant la loterie, tout en permettant des assertions faillibles, telles que «Ma voiture est dans l'entrée».

1. Introduction

Statements about lotteries raise parallel problems for epistemologists who want to articulate conditions for knowledge and those working on norms of

***Dialogue* 52 (2013), 523–549.**
© Canadian Philosophical Association/Association canadienne de philosophie 2013
doi:10.1017/S0012217313000656

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

assertion.[1] In both cases a central concern is whether one can know or warrantedly assert[2], before the draw is made, sentences like:

(A) Ticket n will lose.

A growing consensus has formed that (A) and propositions like it are neither knowable nor assertible. Some writers on the topic, such as Williamson, have argued that this fact is best explained by, and thus provides evidence for, the Knowledge Norm of Assertion (KNA).

KNA: Assert that *p* only if you know that *p*.

Others, such as John Hawthorne, have furthermore argued that ordinary propositions like, "I will not have a heart attack in the next minute," and, "My car is in the driveway" are, properly speaking, lottery propositions.[3] That is, they are relevantly similar to (A) with respect to their knowability and assertibility given probability of less than 1 on one's evidence. An implication of the latter position is that since knowledge is not possible in the traditional lottery case, it is also not possible in the more general cases of everyday life. At least, that is the implication unless a principled line can be drawn between ordinary propositions and traditional lottery propositions.[4]

This paper is composed of two main projects. In the first part, I present the problem for knowledge and assertibility raised by Williamson's and Hawthorne's analyses of the lottery paradox. Williamson's presentation of the argument indicates an inference from the unassertibility of lottery statements, to their unknowability, to merely probabilistic evidence being always insufficient for knowledge. As others have noted, Williamson's position, particularly when combined with Hawthorne's, implies something close to an error theory about knowledge and warranted assertion. I will argue that this consequence is both powerfully counter-intuitive on its own terms and represents a methodological tension for Williamson inasmuch as his analysis uses intuitions about the unknowability and unassertibility of lottery propositions as data.

In the second part, I will offer an analysis of the lottery paradox. This analysis accommodates the intuitions invoked by Williamson, to the effect that one does not know and may not assert (A) in the lottery case. And, unlike Williamson's account,

---

1 Consider Kyburg (1961), Harman (1973), Vogel (1990), Lewis (1996), Williamson (2000), Hawthorne (2004), Kvanvig (2009), among others.

2 I will generally use "assert" to mean "properly assert" or "warrantedly assert" unless otherwise noted.

3 Hawthorne (2004) gives credit for the initial argument to Vogel (1990).

4 Of course, neither Hawthorne nor Williamson are sceptics: they each have an account detailing how we can come to know much of what we intuitively think we know.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

it does not violate equally strong intuitions about knowledge in the ordinary case. I draw on a family of related proposals defined by their sharing the idea that knowledge or assertibility tolerates probabilities less than unity, when the space of possible outcomes includes some that can be properly ignored from consideration.

What it is to *properly* ignore some possibilities when considering defeater conditions for knowledge and assertibility is a point on which different theories may disagree; but, in broadest outline, this family of approaches has the virtue of allowing us to explain an intuitively clear asymmetry between 'lottery' propositions and ordinary propositions (which I will call 'fallible' propositions). I argue that the reason one cannot know or assert a lottery proposition is that one lacks sufficient epistemic justification for such propositions. I finish by drawing some implications for the norms of assertion literature.

2. The Lottery Paradox

Suppose that there is a fair lottery with n tickets, say 1000, where one and only one ticket will win.[5] The probability of a given ticket losing is 99.9%. Suppose that these grounds are sufficient to rationally believe that one's ticket will lose. Jane holds one of these tickets and on these grounds comes to believe *B1*.

B1: My ticket will lose.

By parity of reasoning, if Jane is justified in believing B1, she will be justified in believing *B2*, for any ticket n.

B2: Ticket n will lose.

If we suppose some reasonable formulation of the closure principle, then Jane will be justified in believing *B3*.

B3: All tickets will lose.

However, if this is a fair lottery and one ticket must win, then the probability of B3 is 0 and it's therefore irrational for Jane to believe it. Thus, Jane is both justified and not justified in believing B3, producing a paradox.

This has been a problem for epistemology ever since its first formal presentation by Henry Kyburg in 1961. However, I do not wish to enter into the

[5] In fact, the paradox will result even if it is a biased lottery or even if it's likely that no tickets will win. Williamson (2000, p. 248), DeRose (1996), Douven (2006, p. 462, n.18), and Hawthorne (2004, pp. 8-9). Furthermore, even if it's possible that no ticket will win, Jane has very good evidence that the proposition, "All tickets will lose," is false, such that the paradox still holds.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

discussions of various attempts to resolve the paradox.[6] Instead, I raise it as a foil for the role of the lottery paradox in the literature on norms of assertion. There is an emerging consensus that, prior to information on the results of a draw, one is not warranted in asserting *A1*.

A1: Your ticket will lose.

While there is no consensus on how to best explain this intuition, a prominent attempt belongs to Timothy Williamson. I turn now to his analysis.

3. Williamson and the Lottery Paradox

Williamson begins with the observation that there is a widely shared intuition in our linguistic practices that assertions such as A1 are inappropriate; that is, not warrantedly assertible. Call this intuition *INT1*.[7]

INT1: One may not assert, prior to the release of information on the result of the draw, that a ticket will lose in a lottery.

One's evidence in the case of lottery propositions such as B1/A1 is that it is highly probable that they are true. Suppose that Ben were to make the flat-out assertion A1 to Jane. If Jane were to find out that Ben's evidence for his assertion was merely that it is highly probably true, then even if it happens to be true that Jane's ticket will lose, she is entitled to regard Ben as having spoken improperly. Her taking the view (and perhaps her accompanying resentment!) reflects an intuition that asserting A1 based merely on its being very probably true is tantamount to cheating at the practice of assertion: Ben exceeds his evidential

[6] One might think that any reasonable formulation of the closure principle that gives us B3 is therefore *un*reasonable. The difficult question of whether any closure principle is truly reasonable, all things considered, is not one that I will consider in much detail. However, in my opinion, a promising response is to deny the principle of closure on which an agent justified in holding both conjuncts is thereby justified in holding their conjunction. Kyburg (1961) suggested just such a strategy for resolution. Hill and Schechter (2007) recognize that the aforementioned closure principle will aggregate risk, and that when one conjoins two probabilistic statements, the conjunction will necessarily have a lower probability than both conjuncts. Thus two justified beliefs, which individually meet the standard of justification, may fail to meet that standard when combined. This would dissolve the paradox. Williamson (2000, p. 248) appears to concur.

[7] Williamson is not alone in suggesting that this is a common intuition. See, for example, Hawthorne (2004), Douven (2006), and Kvanvig (2009). I share the intuition that one does not know in the lottery case (INT2), but not that one may never warrantedly assert (INT1). However, for the purposes of this paper, I assume that such assertions are never warranted.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

authority when he asserts A1. The nearest thing to A1 that Ben can legitimately assert is the explicitly qualified *A2*.

A2: It is highly likely that your ticket will lose.

Linguistic evidence for these claims can be adduced from how we consider the response, "But you don't know that!" a decisive criticism of A1 but not of A2.[8]

A prominent feature of the lottery paradox is that the particular probability of the truth of A1 is not at issue. The intuitions about assertibility canvassed above seem to survive on any lottery, no matter how large the number of tickets. So we can make the prior probability that a given ticket will lose arbitrarily close to certainty without undermining the intuition that an agent may not assert A1 (*viz*., INT1) and cannot know B1. Call this intuition *INT2*.

INT2: One does not know that a given ticket will lose, in the absence of information on the result of the draw.

Williamson concludes that merely highly probable evidence, even when the proposition is true, does not suffice for knowledge or warranted assertibility.[9] Call this the *Insufficiency Principle* (*INS*).

INS: That *p* is true and that (one knows that) *p* has (arbitrarily) high probability <1 on one's evidence are insufficient for one knowing that *p*.

This principle generalizes, for Williamson, beyond the case of lotteries. When one's evidence for the truth of a (true) proposition is merely probabilistic (even if it is arbitrarily close to certain), such evidence is always insufficient for assertibility and knowledge.

Williamson's analysis of the lottery paradox differs importantly from the epistemological analysis presented in Section 2. The more traditional epistemological analysis requires some principle of closure; however, Williamson denies that justification is closed (at least in these circumstances). He offers an analogy.[10] Suppose that a large pile of chocolates sits in front of me. It's a large pile, but one that I have a strong desire to completely consume. However, I have been informed that one chocolate has been poisoned. It will make me ill

[8] Williamson (2000, pp. 247-8). He argues that "probabilistic evidence warrants only an assertion that something is probable'" (p. 248).

[9] Later, I will briefly return to discuss Williamson's view of evidence and evidential probability on which knowledge that *p* is probability 1 on one's evidence. I don't mean to here suggest that Williamson arrives at his view of evidential probability *because* of his views of the lottery paradox, though it's possible that this is the case.

[10] Williamson (2000, p. 248).

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

but I will recover. If we suppose that the pile is very large (and I am very hungry), "I can quite reasonably eat just one, since it is almost certain not to be contaminated, even though, for each chocolate, I have a similar reason for eating it, and if I eat all the chocolates, I shall eat the contaminated one, and my sickness will be overdetermined."[11] However, "no plausible principle of universalizability implies that, in the circumstances, any reason for taking one chocolate is a reason for taking them all; the most to be implied is that, in the circumstances, any reason for taking one chocolate is a reason for taking any other chocolate instead."[12]

Williamson thus denies that the justification to take a chocolate, although such justification is present for taking any single chocolate, is closed such that if I am justified in taking one chocolate, then I am justified in taking two chocolates … then I am justified in taking all chocolates, which is absurd.[13] Yet the intuition remains that even though I may be justified in taking a chocolate, I cannot be said to *know* that a given chocolate is uncontaminated. The chocolate case is analogous to the lottery case since, *mutatis mutandis*, INT1 and INT2 apply here also. I may not assert that a given chocolate is uncontaminated (before I eat it) and I do not know that a chocolate is uncontaminated (before I eat it).

How, then, do these various intuitions and principles fit together? Williamson begins with the widely shared intuition that assertions such as A1 are inappropriate (INT1). There is linguistic data supporting this. An intuitively acceptable response to such assertions is to object that the speaker does not *know* that the ticket will lose (INT2). The intuition is that, as far as the speaker knows, the ticket might win; that is, since the assertion is made before information on the outcome of the draw, the speaker does not know that the ticket will lose.

Williamson is a central figure who argues that the constitutive norm of assertion is KNA. He argues that the best explanation of why assertions such as A1 are not warrantedly assertible is that the speaker lacks knowledge. He observes that the speaker's evidence for A1 is merely that it is (arbitrarily) highly probably true; moreover, even when it is true that the ticket will lose, the speaker lacks

---

11 *Loc. Cit*.

12 *Loc. Cit*. Set aside the empirical disanalogies between lottery cases and the diminishing appeal of eating chocolate as one eats more of it. Williamson's analysis of the chocolate-eating case parallels the epistemological analysis of the lottery paradox, on which the justification for believing that any one ticket will lose applies equally to all other tickets, and closure seems to license the move from believing of each ticket that it will lose to the obviously false belief that all tickets will lose. Believing that all tickets will lose is analogous to its being overdetermined that I will become sick from eating the contaminated chocolate.

13 Others have suggested a similar tactic to the lottery case in both the epistemological literature and the assertion literature—see note 5.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

knowledge. So the best explanation for why the speaker lacks knowledge is that merely highly probable evidence, even when true, is *never* sufficient for knowledge.[14] So, Williamson uses INT2 to explain INT1 and then produces the INS principle.

Of course, Williamson's commitment to merely probabilistic evidence being insufficient for knowledge may be independent from his treatment of the lottery paradox. For Williamson, knowledge is evidence. Moreover, all and only knowledge is evidence. Thus, E(vidence) = K(nowledge). From this, he argues, knowledge that *p* is probability 1 on one's total evidence. If E = K, then part of my evidence for my knowledge that *p* is my knowledge that *p.* So even if p(e) = 0.1, p(e|e) = 1. And if I know that *p*, then p(p|e) = 1, since part of my total evidence is my knowledge that *p* (from E=K). That is, when I know that p, that knowledge is added to my set of evidence for *p*. And so we can see how it follows that merely probabilistic evidence is insufficient for knowledge: conditionalizing on merely probabilistic evidence doesn't produce the result that p(p|e) = 1.

This argument raises some concerns, though. One concern is the conclusion that merely highly probable evidence, even when the proposition is true, is never sufficient for knowledge. I will argue that while this is true in the case of lotteries, it's not true for many ordinary propositions such as "I'm holding my coffee cup." Another concern is that this may commit Williamson to a strongly counter-intuitive account of knowledge and warranted assertibility claims and ascriptions. The problem comes from his argument using an intuition as a datum. If using an intuition leads to a strongly counter-intuitive and revisionist account, then this casts doubt on the evidentiary status of the initial intuition that was used as a foundation for his argument. However, I only raise these concerns now to set them aside until Section 5.

4. Hawthorne and Widening the Sense of Lottery Propositions

Following Vogel (1990), John Hawthorne (2004) argues that we should consider a wide class of propositions to be lottery propositions. Lottery propositions share two characteristics: they are (arbitrarily) highly probable on one's evidence and there is an intuitive reluctance to say that one can know or warrantedly assert them. Examples abound. For example, *A3* is a lottery proposition.

> A3: I will not have a heart attack in the next 5 minutes.

Suppose that I am a healthy adult and I have just finished undertaking a rigorous physical exam by the best team of doctors in the world. They assure me that my health is excellent. In these epistemic circumstances, we may fairly describe the probability of my suffering a heart attack in the next 5 minutes as

[14] Williamson (2000, p. 250).

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

exceptionally low. Yet Hawthorne argues that we can reconstruct the case as my holding a lottery ticket where, oddly enough, 'winning' is having the heart attack. Since it is very unlikely that I will 'win,' the lottery has many tickets. Let us suppose that there is a one in a million chance of my suffering the heart attack. But even if the lottery has one million tickets, Hawthorne argues that the intuitions INT1/INT2 remain even if it is true that I will not suffer the heart attack. For, INT1/INT2, by Williamson's argument, arise even if the truth of the proposition is arbitrarily close to certain and even when it is true. Thus, Hawthorne argues that if a proposition can be properly reconstructed to be the outcome of a lottery, then one can neither know nor warrantedly assert the proposition.[15]

The key to Hawthorne's proposal is parity reasoning. In the original lottery case, Jane doesn't have any more reason for thinking that her ticket will lose than she has for thinking that any other ticket will lose. So, at the root of lottery propositions is that they can be conceptualized as being divided into a "possibility space of a set of subcases, each of which, from the point of view of the subject, is overwhelmingly likely to not obtain, but which are such that the subject's grounds for thinking that any one of the subcases does not obtain is not appreciably different than his grounds for thinking that any other subcase does not obtain."[16] He goes on to write,

> In general, what is often at the root of the relevant lottery intuition is a division of epistemic space into a set of subcases with respect to which one's epistemic position seems roughly similar. Once such a division is effected, a parity of reasoning argument can kick in against the suggestion that one knows that a particular subcase does not obtain, namely: If one can know that *that* subcase does not obtain, one can know of each subcase that it does not obtain. But it is absurd to suppose that one can know of each subcase that it does not obtain.[17]

Hawthorne thus proposes the following principle of *Parity Reasoning*:

> 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 of the set will not obtain than one has for thinking that any other particular member will not obtain. Insofar as one reckons it absurd to suppose that one is able to know

[15] Hawthorne attempts to argue for a principled distinction between cases that can, in principle, be reconstructed as lotteries but where knowledge is nevertheless possible. It's a 'salience' version of a relevant alternatives theory (RAT). I discuss such attempts beginning in Section 5.

[16] Hawthorne (2004, pp. 14-15).

[17] Hawthorne (2004, p. 15).

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

> of each (p_1, … p_n) that it will not obtain, one then reckons oneself unable to know that p.[18]

We can thus see how a proposition such as, "I will not have a heart attack in the next 5 minutes," can be reconstructed as a lottery proposition such that Parity Reasoning can kick in. Call this proposition *H*. Divide up the possibility space for my having a heart attack in the next 5 minutes. Suppose that the probability of my having a heart attack is vanishingly small, but nonzero. Relative to my epistemic position, my reasons for thinking that, in this case, H is true are not appreciably different than in cases where H is false. Therefore, by Parity Reasoning, if I can't know of each and every subcase that H is true, I cannot know in this subcase that H is true. The upshot of Hawthorne's argument is that a proposition can be reconstructed as a lottery proposition when the proposition can be conceptualized where the possibility space can be divided such that Parity Reasoning applies.

It's natural to think that lottery propositions only cover instances of prediction, since there's an inherent uncertainty about the future. However, another of Hawthorne's examples suggests that many propositions about the present are also lottery propositions. Consider *A4*.

> A4: My car is parked in the driveway.

Suppose that I have just looked to my driveway and clearly saw my car through my front window. My car really is in the driveway. Normally, we take this to be sufficient to count as something that I know. But suppose that before I make the assertion A4, I turn around so that I no longer see my car. Do I still know? It's at least *possible* for my assertion to be false: given my evidence about the location of my car just a moment ago, it's possible for it to have been stolen. Of course, given my background understanding of how long it takes for someone to steal my car, I consider the possibility extremely remote (even nearly impossible). And we can reconceptualize the proposition into a possibility space where the proposition is true in the vast majority of subcases, but my reasons for thinking that A4 is true in this subcase are not appreciably different from subcases where A4 is false. So, pace Hawthorne, this is still a lottery proposition and INT1/INT2 should remain even if it's the case that my car really is in my driveway (*viz*., the proposition in A4 is true). Therefore, unless a distinction can be made to save the knowability of propositions like A4, Hawthorne's argument implies that we cannot know many of the propositions that we ordinarily take ourselves to know.[19]

---

[18] Hawthorne (2004), p. 16.

[19] Hawthorne eventually proposes a salience version of a relevant alternative theory in order to save our ordinary claims to knowledge.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

While the car example demonstrates that lottery propositions range over claims about the present, the pervasiveness of lottery propositions extends even into statements about the past. Consider *A5*.

A5: The Magna Carta was signed in 1215. [20]

We may have very strong evidence for the truth of the proposition in A5, such as explanatorily coherent and mutually confirming separate accounts that indicate the year as 1215. Still the possibility remains that all of these accounts are mistaken. So, Hawthorne would suggest, we could consider ourselves as holding a ticket to the *The-Magna-Carta-was-not-signed-in-1215* lottery. We expect to lose since it is unlikely that the Magna Carta was not signed in 1215, but it is possible that it was not. And we can reconceptualize the proposition such that Parity Reasoning would apply. Since it's possible that we could 'win' this lottery, and our reasons for thinking that A4 obtains in this subcase are not appreciably different from cases where A4 is false, Hawthorne would argue, INT1 and INT2 suggest that we do not know and may not assert A4.

Lottery propositions therefore range over instances of predictions, retrodictions, statements about the present, and perhaps others. They pervade the sorts of propositions we intuitively take ourselves to know and warrantedly assert. I may have very good evidence that I'm holding my favourite coffee cup at the moment, but if I close my eyes and assert, "I am holding my coffee cup," I must admit that it's possible that someone has quickly and surreptitiously switched my cup for something that has the identical weight and feel of my cup. Call this assertion *A6*.

A6: I am holding my coffee cup.

Since, considering the possibility space, my reasons for thinking A6 true are not appreciably different in subcases where A6 is false, A6 is a lottery proposition.

Hawthorne's analysis of the generality of the lottery paradox is an extension of Williamson's analysis. Hawthorne argues that propositions that can be reconstructed as lotteries, and where Parity Reasoning applies, are lottery propositions. The important implication is that when we can properly perform

[20] One might think that there's a difference between this case and the previous cases. The former all seem to be both epistemically possible and physically possible, whereas the latter is merely epistemically possible (if we assume that the Magna Carta was signed in 1215). For epistemologists such as Williamson and Hawthorne, these cases only need to be epistemically possible to produce the relevant problem. See Hawthorne (2004, pp. 111-2, and elsewhere) where he discusses the Epistemic Possibility Constraint.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

this reconstruction, one's evidence is merely probabilistic in Williamson's sense. Thus, insofar as it is not the subject matter *qua* lotteries that makes such propositions unknowable (and not warrantedly assertible) but rather that one's evidence is merely probabilistic, Williamson's position appears to lead to Hawthorne's.

A striking implication of these arguments is that we tend to be radically mistaken about what is knowable and assertible. This would suggest something close to an error theory or scepticism about knowledge and warranted assertibility claims and ascriptions.[21] We are just generally mistaken when we think that we know or may assert something. I'm confident in my intuition that I know and may warrantedly assert that I am holding my coffee cup, even immediately after I close my eyes. I have good reasons to think so, despite recognizing the possibility that I am mistaken and therefore recognizing that I could reconstruct the case as a lottery where Parity Reasoning would apply. The arguments of Hawthorne (and Williamson) suggest that when my intuitions conflict with the results of this analysis of lottery propositions, then so much the worse for my intuitions. This strong counter-intuitiveness is a problem shared by Williamson and Hawthorne.

5. An Alternative Analysis of the Lottery Paradox

Let's assume, for the sake of argument, that in the lottery case (B1/A1) one does not know despite having probabilistic evidence that indicates that the proposition is (arbitrarily) likely to be true. It is reasonable to conclude that having a highly probable, true belief is not always sufficient evidence for knowledge. However, it is an entirely different matter to think that having a highly probable true belief is never sufficient for knowledge. Since nearly any proposition can be reconceptualized as a lottery proposition (or one that entails a lottery proposition), the suggestion that highly probable evidence is never sufficient for knowledge is what opens the door to the error-theoretic implications of Williamson's and Hawthorne's arguments. Thus, even our ordinary knowledge claims seem based on merely highly probable evidence. It remains to be seen whether a principled line can be drawn between 'ordinary' and 'lottery' cases in order to save the intuition that we often do know in the ordinary case (an intuition both Williamson and Hawthorne want to retain). One strategy is to establish that in some class of cases merely highly probable evidence for true propositions is sufficient for knowledge provided that the agent has the right sort of merely probabilistic evidence. I turn now to provide such a proposal.

First, I suggest making a distinction between what have come to be called "lottery propositions" in the literature and what I will call "fallible propositions." I suggest that we reserve the name "lottery propositions" for propositions about

[21] Douven (2006) and Dodd (2007), among others, have also noticed this potential consequence of Williamson's position.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

genuine lotteries. In this sense, A1 is a lottery proposition but A6 is not. Let us call "fallible propositions" those propositions that Hawthorne argues may be reconstructed as a lottery but are not actually about lotteries. In this sense, A3-6 are fallible propositions.[22] This distinction will facilitate establishing a principled difference between the two kinds of propositions, pace Hawthorne.

A promising strategy for providing a principled distinction between lottery and fallible propositions can be found in a family of related proposals sharing the idea that knowledge or assertibility tolerates probabilities less than unity, when the space of possible outcomes includes some that can be properly ignored; that is, properly excluded from consideration. What it is to *properly* ignore some possibilities when considering defeater conditions for knowledge and assertibility is a point on which different theories may disagree; but in broadest outline this family of approaches has the virtue of allowing us to explain an intuitively clear asymmetry between lottery and fallible propositions. I will refer to them as "relevant alternative" theories (RATs).[23]

An established example of a RAT is David Lewis' version of contextualism. That it is a contextualist account is not important; there are also non-contextualist RATs, after all. Lewis appeals to possible worlds in explicating the concept of alternative possibilities, and proposes that "*S knows* that *P* iff S's evidence eliminates every possibility in which not-*P*—Psst!—except for those possibilities that we are properly ignoring."[24] The concepts of *eliminating* and *ignoring* alternative possibilities, and not the specifically Lewisian view of possible worlds, are what make this a clear case of a RAT.[25] More generally: in

---

22 One might think that, "I will not be dealt a straight flush on the next hand," is thus a fallible proposition rather than a lottery proposition. I don't think that this is correct; rather, it's a lottery proposition. The probabilities of being dealt a straight flush in a hand of poker are calculable *a priori* (assuming that the event would be stochastic rather than deterministic). That this is true is how casinos manage consistently to earn profits. On the analogy between actions such as shooting a free-throw in basketball and poker (*qua* stochastic events), see McKinnon (2013b).

23 I use this term with some reluctance because there is already a distinction in the literature between specific theories known as, for example, *the* Relevant Alternative Theory (and the "New" Relevant Alternatives Theory), and Lewis' contextualist account as well as non-contextualist accounts. I think that the family resemblance is so strong that it is reasonable to refer to the whole group as relevant alternative theories since they all focus on the feature of dividing the possibility space into relevant and irrelevant alternatives. See Vogel (1999), for example.

24 Lewis (1996, p. 378).

25 For a good discussion, see Hendricks (2006). I do not wish to wade into the debate over what fixes the context; for example, whether it's speaker- or ascriber-dependent. This is a notoriously difficult topic that I believe does not have a significant bearing on my use of the contextualist account of knowledge.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

order for Subject S to know that p, S must eliminate all of the relevant alternative possibilities where $\neg p$ while ignoring irrelevant alternative possibilities where $\neg p$.[26] Call this the *RAT Principle*.

> RAT Principle: In order for Subject S to know that p, S must eliminate all of the relevant alternative possibilities where $\neg p$ while ignoring irrelevant alternative possibilities where $\neg p$.

This is a principle common to all relevant alternative accounts. In order to eliminate a relevant possibility, $\neg p$ an agent will need evidence for $\neg\neg p$. For example, suppose that I see a yellowish bird in a garden and a friend asks me whether there's a canary or a raven in the garden.[27] I respond, "A canary," because, given the only two relevant possibilities for the species of the bird, and since it is yellow and not black (the colour of a raven), I have eliminated the possible world where $\neg p$ *viz*. the raven-in-the-garden world.[28] However, if the question were, "Is there a canary or a goldfinch in the garden?" then I would not have sufficient evidence to assert that there is a canary since I do not have evidence that it is not a goldfinch. In this case, I do not have sufficient evidence to eliminate the alternative possibility where $\neg p$ *viz*. the goldfinch-in-the-garden world.

Providing a precise definition of what constitutes relevant alternatives requiring elimination, or irrelevant alternatives that one is permitted to ignore, is not an easy matter. This is largely where disagreement between various RATs is found. Lewis himself fails to provide a rigorous technical definition (as the self-conscious informality of his formulation indicates!). What constitutes a possibility that one can properly ignore? This is left unclear. Some options are: significance, probability, relevance, resemblance, salience, and pragmatic effects.[29] This list is not exhaustive. Each has its own problems as a precise solution. For example, there is strong reason to doubt that just probability will work. The lottery paradox seems to show precisely that an alternative possibility's being extremely unlikely is not sufficient in isolation to justify ignoring it. Although Jane may be arbitrarily close to certain that she will lose the lottery, she may not properly ignore the alternative possibility where she happens to win.

---

[26] Lewis uses the language (and concept) of possible worlds in place of "alternative possibilities."

[27] An example adopted from Austin (1946) also discussed in Schaffer (2008).

[28] Many contextualists consider knowledge to be infallible within the context. That is, in order for a relevant possible world where $\neg p$ to be eliminated, it must be known that $\neg p$ with certainty. That is, "knowledge in light of possibilities of error sounds like a contradiction" Hendricks (2006, p. 67). However, for my purposes I will allow for the "evidence" that $\neg p$ to be fallible.

[29] Dougherty and Rysiew (2009) offer a "significance" version; Hawthorne (2004) offers a "salience" version; Fantl and McGrath (2009) offer a "pragmatic effects" version.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

Traditional attempts to specify the class of legitimately ignorable alternatives based on significance, relevance, world-resemblance, and salience share a common problem: namely, identifying just what determines whether an alternative possibility counts as relevant, salient, significant, and so on.[30] As a basic necessary condition, Gail Stine (1976) suggests that an alternative is relevant only if one has some reason to think that it's true. This seems intuitive: I know that that's a zebra because I don't have any reason to think that it's just a cleverly painted mule.[31] In fact, I have positive reasons to doubt that it's a cleverly painted mule because, perhaps, the zoo I'm attending is a reputable zoo and wouldn't risk its reputation. However, is the mere possibility of something sufficient reason to think that it's true? This is left unclear. If something is merely possible, or where one's evidence for $\neg p$ in subcases where $\neg p$ is not appreciably different in cases where ¬p is false, is sufficient reason to think that something is true, then this suggests, pace Hawthorne, that all possibilities are relevant, and we arrive at the sceptical problem.

One attendant problem for such views (*i.e.*, salience, significance, etc.) is whether (mere) reflection or idle inquiry can too easily change the set of relevant alternatives that one must eliminate simply by bringing up a previously unconsidered further alternative. Suppose that I am asked whether there is a canary or raven in the garden. In this context, I have evidence against the raven-in-the-garden world and (assuming that there is a canary in the garden) I may be said to know that there is a canary in the garden. But suppose that my interlocutor asks if it's possible that I'm dreaming. On some measure of relevance, this would make the I'm-dreaming world relevant to whether I know that there is a canary in the garden; I would now have to provide evidence against this possibility in order to eliminate it. If I could not, then I would not know that there is a canary in the garden, as I would not have eliminated all of the relevant alternative possibilities. But this seems to unduly privilege the power of an interlocutor to make relevant some unconsidered alternative. Moreover, this would suggest that cautious and reflective epistemic agents, who consider more alternative (and potentially less probable) possibilities, will know less than an incautious agent who considers fewer. The philosopher, after all, may doubt that she has hands, while a carpenter will not. Does the carpenter know that she has hands when the philosopher does not? We would not want this to be a consequence of a RAT.

It is a difficult task to explicate the conditions under which we can properly ignore alternative possibilities by deeming them irrelevant. Difficult tasks beset many theories, so I don't mean to suggest that we should rule out the possibility of an adequate account of what constitutes the necessary and

---

30 This is where I think Fantl and McGrath's pragmatic "idleness" view has an advantage. In fact, I think that all destabilizing alternatives are idle. However, it's not clear whether all idle alternatives are destabilizing.

31 A reference to Dretske (1971).

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

sufficient conditions for properly ignoring alternative possibilities. At a minimum, each member of the RAT family is latching onto something that appears intuitively correct: namely, that there is *some* story worth telling about how one can properly ignore some alternative possibilities in coming to know a proposition. This is especially true if one holds both that one cannot know in the lottery case and that one can know in the ordinary case.

For my part, I suggest that an alternative possibility may be deemed irrelevant if its truth would rationally require an agent to re-evaluate or restructure some of, or enough of, a particular class of background beliefs and/or theories about the world.[32] That is, the supposition of an alternative possibility would or ought to trigger a rational re-examination and restructuring of a significant class of background beliefs about the way the world works in a web-of-belief sense of 'background' beliefs.[33] I will refer to this as destabilization. For my present purposes, I will leave this concept deliberately vague. I will refer to this account as a *Non-Destabilizing Alternatives Theory*, or NDAT.

> NDAT: An agent may ignore an alternate possibility iff the supposition of the alternative possibility would or ought to trigger a rational re-examination and restructuring of a significant class of background beliefs.

The impact on background beliefs in NDAT is qualitative, rather than quantitative. We must avoid cashing out destabilization in terms of the number of background beliefs impacted. Suppose that Jane believes B1 and that she finds out that she actually won the lottery. It is natural to suppose that she would change a tremendous number of beliefs about her future (based on what she can do with the prize). But this is not what I mean. Instead, destabilization refers to the change to more general beliefs and theories about how the world works.

Ignoring possibilities based on their impact on background beliefs is not an entirely unfamiliar concept. There is a general theory of objects such that things such as cups cannot instantaneously (and for no apparent reason) become other

---

[32] Smith (2010) proposes something similar, but importantly different, from my view. He argues that an alternative is relevant only if one's evidence *normically* supports the proposition. Evidence normically supports a proposition when the presence of the evidence but falsity of the proposition demands explanation; evidence that doesn't normically support a proposition doesn't have this property. Losing a lottery doesn't demand a special explanation, so my probabilistic evidence that my ticket will win does not normically support my losing the lottery. However, my seeing my cup and feeling its weight right before I close my eyes normically supports the proposition that I'm holding my cup. If I were to have those experiences and the proposition were false, this would demand a special explanation.

[33] As in, for example, Quine (1978).

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

objects like black holes. Science is concerned with articulating and compiling evidence for such theories. Suppose that Jane forms the belief that her cup will not suddenly become a black hole. Call this *B4*.

B4 My cup will not suddenly become a black hole.

Suppose that it's possible for B4 to be false. The falsity of B4 is, therefore, an alternative possibility. On my view, if Jane were to find out that B4 is false, then this would be a destabilizing event. Her general theory of objects and nomological understanding of the world would or ought to be radically altered. But it's not the *number* as much as the *quality* of the background beliefs that would need to change on the supposition that B4 is false. My argument boils down to this: one may properly ignore destabilizing alternative possibilities in coming to know or assert some proposition.

In addition to involving a certain quality of impact to background beliefs rather than quantity, destabilization is a graded rather than a threshold concept. The supposition (or the coming to be) of some alternative possibilities will be more or less destabilizing than others. If I were to learn that cups could regularly suddenly become black holes, then this would be especially destabilizing to my background beliefs and theories about the world. However, if I were to find that my belief that my car is still where I parked it is false (since it has actually just been stolen), then this may be slightly destabilizing for my beliefs about the security of the faculty parking lot, but it would be far less destabilizing than learning about cups becoming black holes. In the former, I would merely be rationally required to re-examine my beliefs about how quickly thieves can steal a car; in the latter, I would be rationally required to re-examine my views on the basic physical laws and probability assessments of whether objects can turn into singularities.

A useful concept in explicating destabilization is *consilience*. Generally speaking, consilience is a property of a body of beliefs, theories, or evidence, arising in distinct explanatory domains, which are mutually evidentially explanatory. In the philosophy of science, consilience has often been taken as an indicator of good theories, broadly on the grounds that cross-domain, intertheoretic explanatory unity is a powerful element or indicator of correctness (on pretty much any notion of correctness).[34] For example, consider what has

[34] Whewell (1840) was the first to introduce the concept. It has since been used in a number of contexts in the Philosophy of Science. However, I wish to distance my use of the concept from that of Wilson (1998). Consilience is sometimes manifested in a theory's producing successful predictions and explanations in areas where it was not expected to apply. For example, Newton's theory of gravity was not only able to explain the movement of celestial bodies, but also terrestrial motion. Moreover, it also both predicted and explained tidal motion through the gravitational influence of the moon.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

become known as the "modern synthesis" of (evolutionary) biology into what has also been referred to as "neo-Darwinism."[35] The synthesis of Darwin's theory of evolution by natural selection depended heavily on finding a vehicle for genetic heredity. When it was found (in the form of DNA), this led to successes in disparate fields such as economics (evolutionary game theory), psychology (evolutionary psychology), computer science (evolutionary programming and genetic algorithms), and many others. With each new success came more support for the synthesized theories of genetic heredity and evolution through natural selection.

Now, since consilience pertains to hypotheses, theories, and beliefs, even highly consilient theories are consistent with the possibility of disconfirming observations. Some such observations would necessitate small modifications to a body of beliefs or hypotheses, but others would require a more radical restructuring of one's beliefs about *how the world works*. It is these latter events that I am describing as destabilizing. Irrespective of whether we ultimately choose to use the particular notion of consilience to underwrite the concept of destabilization, I think that it's plausible to see the intuitively appealing shared component of RATs as destabilization.[36]

Destabilizing events, though possible, are often ignored as being not worth consideration. I have not given a general definition of what counts as knowledge, though. I have merely been arguing for an account of when we can ignore (irrelevant) alternative possibilities in coming to know some proposition. Some might wonder what it means to knowledge if some destabilizing event, which was ignored as an alternative possibility, actually came about. If subject S ignores $\neg p$ because it is destabilizing, it does not follow that S knows $\neg p$ (this is a mistake that some RAT theorists have made).

Suppose that I'm looking at what I take to be a zebra. I don't have any reason to think that it's a cleverly painted mule, and, in fact, I have good reason to doubt that it's a cleverly painted mule. If it were to turn out that what I'm looking at *is* a cleverly painted mule, then I would consider this destabilizing to my background beliefs about the zoo and about the passability of a cleverly painted mule for a genuine zebra. So I ignore the irrelevant possibility that what I'm looking at is a cleverly pained mule. However, I have only *ignored* the alternative possibility, I haven't *eliminated* it. When an agent eliminates an alternative possibility $\neg p$ (e.g., the raven possibility), the agent comes to know $\neg p$ (i.e., that the bird in my garden is not a raven). However, when an agent merely ignores an alternative possibility $\neg p$, the agent does not come to know $\neg p$. Ignored alternatives do not thereby become known. So in ignoring the irrelevant possibility that what I'm looking at is a cleverly painted mule, I don't thereby know that it's not a cleverly painted mule.

---

[35] For a detailed description, see Mayr (1991).

[36] This is where my view dovetails with the view expressed in Smith (2010).

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

Now, one might question whether these alternatives are really destabilizing. One might think that, in the zoo case, it's more than merely possible for some pranksters to sneak in and cleverly paint a mule to look like a zebra. Moreover, one might think that in the parked car case, although I just saw my car parked in my driveway seconds ago, maybe it *was* just stolen: cars are stolen all the time, after all.[37] But just as highly consilient scientific theories may consider some event highly unlikely, sometimes destabilizing events *do* occur and theories are discarded or radically altered. If I were to find out that the zebra-like animial that I'm looking upon at the San Diego Zoo—a zoo with very good security and a very good reputation—is really a cleverly painted mule, and the zoo didn't know about it, this would cause me to seriously reconsider my beliefs about the zoo, the probability that pranksters could pull off such a prank, the reliability of my eyesight (perhaps in picking out zebras), and a host of other beliefs that would be impacted by this outcome. This is what makes it destabilizing. The same would be true, *mutatis mutandis*, if I were to find out that I was wrong about my car being in the driveway mere moments after my having looked at it. Of course, my various background beliefs in these contexts don't rule out my being wrong, but that's a common feature of destabilizing events.[38] And certainly, finding out about my being wrong about the car being stolen would be less destabilizing than the case about a coffee cup becoming a black hole, but that goes to my point about destabilization being a graded concept.

The interplay between background beliefs, ongoing belief-formation, and observations of the world is an important feature of ordinary non-lottery cases such as A6 and fallible propositions in general. If John were to find that, when he asserted A6, someone had actually surreptitiously replaced his cup with a similar feeling object, this would create a rational requirement to reform his background beliefs about the feasibility of making such a surreptitious switch quickly and without his noticing.[39] In this case, if John were to find out that his belief about holding his coffee cup were false (because, in the instant between perceiving his cup and closing his eyes, someone replaced his

---

[37] I thank an anonymous referee for this objection.

[38] Sometimes, though, destabilizing events may be (in some sense) ruled out by background beliefs, and these cases rationally require an even more radical re-structuring of background beliefs. However, this would require a specification of 'ruling out' some possibility, which is beyond the scope of this paper. I also leave 'rationally required' somewhat unspecified.

[39] One might wonder whether NDAT suggests that one may properly ignore some relevant alternative, based on one's background beliefs, even if the alternative is believed to be true by other epistemic agents (and we can assume that they're epistemic peers). This is a difficult question that would require wading into the epistemology of disagreement, and so it's outside the scope of this paper.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

cup with a cup-feeling non-cup), then this would be a destabilizing event. However, he wouldn't think that it was sufficiently analogous to the lottery case as though he just won the lottery in the I'm-not-holding-my-coffee-cup lottery.

This is consonant with the sense that there is something importantly different between lottery propositions like A1 and fallible propositions like A6. In ordinary cases our background beliefs and beliefs about underlying mechanisms are, or ought to be, importantly sensitive to outcomes. One of the reasons why lottery-talk is so useful, by contrast, is that it abstracts away from questions about fairness, randomness, or underlying mechanisms. No matter how unlikely winning a given lottery may be, the prospect of winning requires no change to one's background beliefs about how the world works. Indeed, if Jane were to win her lottery and thereby reform her beliefs about the likelihood of winning a similar lottery, we would criticize her for committing a statistical fallacy. Suppose that Jane wins her 10,000 ticket lottery (and it is a fair lottery) and, on the basis of this, forms the belief that the odds of her winning must have been greater than 1:10,000. It would be appropriate to chastise her for committing a fallacy. Winning a large lottery may be unexpected, but someone must win, so, why not her? Of course, winning an unlikely event may be evidence for an alternative explanation (such as a rigged draw), but in the lottery case, this is insufficient to justify the belief. Thus, there seems to be something importantly different between the lottery and ordinary cases.

Now, one might object that this seeming distinction can be blurred by consideration of a brain-in-a-vat (BIV) hypothesis.[40] For example, suppose that Bill is a BIV. The world is generally exactly the way that Bill thinks it is except that he is the result of an ambitious neuroscience experiment: he's the first viable BIV ever! The objection is that Bill, upon learning of his BIV status (suppose that he could come to know this somehow), wouldn't feel a rational requirement to change his more general background beliefs about how the world works.[41] For example, he would still be able to make predictions about people's behaviour and the interaction of objects as well as he could before. He

[40] My thanks to John Turri for raising this objection.

[41] There's an interesting question that might be raised here: normally, BIV hypotheses are the sort that are difficult, if not impossible, to come to know. So wouldn't it be the case that difficult-to-know, or downright unknowable propositions, couldn't be deemed irrelevant? On the contrary: it would seem that such alternatives are even better candidates for irrelevance, since they tend to be particularly destabilizing and outside the possibility (or likelihood) of elimination. Our only option is to ignore them (or else give in to the sceptic). But if we assume that we don't need to (or want to) concede the sceptic's point, how should we go about justifying ignoring such possibilities? That's exactly the purpose of NDAT. My thanks to an anonymous referee for this objection.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

would just have to repair his local beliefs about his relationship to the world and people in it *qua* his newfound knowledge of being a BIV. The only difference his discovery makes is that being a BIV is now a live alternative possibility. Therefore, Bill's learning that he is a BIV may not be a clear case of a destabilizing event. Since I have argued that we may ignore destabilizing alternative possibilities, this sort of BIV hypothesis would not be a proposition that Bill (or anyone else) could reasonably ignore as an irrelevant possibility prior to learning that he is a BIV. This would be contrary to the intuition that BIV hypotheses are irrelevant alternative possibilities. Therefore, the objection is that NDAT cannot explain why we ignore BIV-style alternatives, since such possibilities may not be destabilizing.

The objection is not compelling, however. It asks us to consider Bill's epistemic state upon learning that he is a BIV even though the world is generally how he thinks that it is. If, prior to learning that he is a BIV, Bill thinks that there are tables and chairs, that the Toronto Maple Leafs are unlikely to win the Stanley Cup this year, and that it snows in Canada, then all of these propositions are still true. However, his beliefs about his relationship to the world are all false. If this objection were successful, then we would expect that learning that he is a BIV, even though the world is generally the way that he thinks it is, would not trigger a restructuring of Bill's background beliefs about the world. Bill could conceivably still go about placing wagers on the Maple Leafs not winning the Stanley Cup this year and do so rationally. After all, his beliefs about the quality of the team are just as true as they were before he learned about being a BIV. All that has changed is that he no longer actually has hands or the physical relationship to the world that he thought he had prior to learning about being a BIV.

Since I have argued that the impact on beliefs is qualitative rather than quantitative, perhaps such a case would not be as destabilizing as traditional BIV hypotheses (where the world is not generally as we think it is).[42] I have argued that destabilization is, after all, a graded concept: some events may be more or less destabilizing than others. I suspect that the intuition surrounding learning that one is a genuine BIV is that such circumstances would be strongly destabilizing; though, admittedly, perhaps less so than more traditional BIV/evil-demon hypotheses.[43] The sorts of propositions that one forms are heavily influenced by one's beliefs about one's relationship to the world. If Bill were to learn that

[42] One might ask how my view differs from a variety of possible worlds accounts where we can safely ignore 'distant' possible worlds. In a sense, my view is a version of such accounts, and how we specify 'distance' is in terms of destabilization. The more destabilizing an event is, the more 'distant' the possible world.

[43] I say 'genuine' since we're meant to imagine Bill being a brain-in-a-vat rather than something less amazing such as learning that brains-in-bodies are, in some sense, brains-in-vats.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

he is a BIV, then it's hard to predict how he would react, but our popular culture apparently shares the intuition that he wouldn't take it very well: it would be highly destabilizing.[44] Bill's preferences for watching the Maple Leafs each year may be predicated on the benefits of winning wagers, such as being able to afford better tickets for the next game. But if he is a BIV, then there's no benefit to winning: the neuroscientists could merely provide a different input to give Bill the sensation of being in the best seats in the arena. It's plausible to consider many of one's life plans as strongly connected to one's beliefs about one's connection to the world. So, learning that one is a BIV, even if the world is generally the way that one thinks it is, would be highly destabilizing. I suspect that one would change many of one's life plans upon learning that one is a BIV.

One might also wonder why we should take a rational requirement to re-evaluate one's background beliefs as good reason to ignore an alternative possibility.[45] The starting point here is a strong intuition that there's an important epistemic difference between lottery and fallible propositions such that we can know the latter but not the former. Naturally, one would hope to find an epistemic feature that explains this difference. How evidence and alternate possibilities affects our beliefs is one place in which to look, and that's what NDAT does: it locates a difference in the epistemic status of fallible and lottery propositions in an epistemic feature: destabilization.

6. Re-Examining Hawthorne

This notion of an asymmetry between lottery propositions and fallible propositions, along the lines of destabilization, can be more fully elucidated by discussing one of Hawthorne's examples (taken from Vogel): the Heartbreaker hole-in-one case.[46] Sixty (amateur) golfers in a tournament must all play a short, but difficult, par three known as the 'Heartbreaker' hole. Holes-in-one in golf are very rare, even for touring professionals. For professionals the odds of a hole-in-one are roughly 1 in 3000; for an amateur roughly 1 in 12,000.[47] Based on this, prior to the tournament, it would be reasonable to form the belief *B5*.

B5: Not all 60 golfers will get a hole-in-one on the Heartbreaker.

Roughly, the odds of all 60 golfers getting a hole-in-one on the Heartbreaker (*i.e.*, B5 being false) are astronomical: 1 in 5.63×10^{244}. This number is larger than all of the nanoseconds in the universe since the Big Bang by many orders

[44] Consider movies such as *The Matrix* (1999), *Possible Worlds* (2000), and *Vanilla Sky* (2001).

[45] I thank an anonymous referee for raising this objection.

[46] Hawthorne (2004, pp. 12-13 and 17); Vogel (1999, p. 165).

[47] Kindred (1999). I have rounded down for simplicity.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

of magnitude. Hawthorne argues that the appropriate response to finding oneself in the situation where all 60 golfers make a hole-in-one is the same as finding oneself in the situation where one wins an extremely large lottery. This would suggest that no restructuring of one's background beliefs would be required because no such response to winning a lottery is required (or perhaps even permitted, because it would be irrational). However, I think that this does not accurately capture our intuitive responses to such cases as the Heartbreaker.

Consider what it would really be like to witness the event of all 60 golfers making a hole-in-one. The intuition, I think, seems to be that after, say, eight or nine holes-in-one, the next hole-in-one would cease to be surprising. Why? Partly because such an event is astronomically unlikely, we'd search for a more probable explanation; namely, that the game is rigged to result in holes-in-one on the Heartbreaker. One of the structural presuppositions of being able to interpret the ordinary cases as lotteries is that they're fair—there's no foul play. Additionally, the possibility space must be divided such that one doesn't have any appreciably better reasons for thinking that one subcase obtains rather than another. However, partly based on background beliefs about the implausibility of all 60 golfers honestly making holes-in-one, we ignore this possibility and search for alternative explanations, such as cheating, or a local topography that funnels shots into the cup. Even if no such better explanations were immediately forthcoming, the intuition would still be to avoid believing that all 60 golfers achieved their feats under the sort of conditions we initially took to obtain when we first evaluated B5. We intuitively ignore that possibility and assume that some conditions or other obtain under which the supposition of the falsity of B5 would not be destabilizing. We would first attempt to exhaust all better explanations before being driven to the belief that each golfer honestly makes a hole-in-one.

However, if we suppose that the golfers are honestly making these holes-in-one, I suggest, we would intuitively be inclined to re-examine our background beliefs about the likelihood of such an event occurring. This would be destabilizing. However, contra Hawthorne, if the Heartbreaker case can be properly reconstructed as a lottery, then such an inclination would be absent. The presence of the inclination to re-examine our background beliefs and theories about the world (given that all 60 golfers honestly made their holes-in-one) suggests that the Heartbreaker case cannot be properly reconstructed as a (very large) lottery. There seems to be an intuition that one would not be committing a statistical fallacy by re-examining the likelihood of the Heartbreaker event given its occurrence (because we re-examine our background beliefs about the world). Therefore, this indicates an important asymmetry between lottery and fallible propositions. In fact, if Jane were to repeatedly win large lotteries, this would constitute stronger evidence of cheating (much like, I have argued, the Heartbreaker case) and a stronger candidate for destabilization.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

For Hawthorne, mere attention to the possibility of error is not sufficient to vitiate knowledge. Instead, knowledge is vitiated by attention to an epistemic possibility of error. He calls this the *Epistemic Possibility Constraint* (EPC).[48]

EPC: If it is epistemically possible for S that $\neg p$ then S does not know that p.

However, in order for us to take the sceptic seriously, the sceptical alternative possibilities of fallible propositions must be epistemically possible. If John is holding his cup and is asked whether he can know A6 because, after all, it's possible that his cup has been surreptitiously switched, John will entertain such an alternative possibility as epistemically possible. In fact, it's not clear what difference Hawthorne has in mind between the mere attending to a possibility and something being epistemically possible. Typically, something is epistemically possible for an agent if some proposition (or state of affairs) is consistent with what one knows.[49] All of the alternative possibilities discussed count as epistemically possible in this sense: whether one is a brain-in-a-vat, or whether a coffee cup has been surreptitiously replaced with a non-cup, whether a table has temporarily changed due to a massive confluence of quantum effects, or whether all 60 golfers honestly make a hole-in-one on the Heartbreaker.

As indicated above, Hawthorne argues that, if it is epistemically possible that $\neg p$, then one cannot know that p (EPC). But this is just what it means to be a fallibilist: that one may know p despite $\neg p$ being consistent with one's evidence.[50] Fortunately, since most people favour fallibilism, we can simply reject EPC and I have given grounds for doing so. Consider John entertaining the possibility that A6 is false because someone has switched his cup. Naturally, he will ignore the possibility. He recognizes and accepts it as an epistemic possibility, but he dismisses it as irrelevant. If A6 is true, and he has good reason to think that it is based on his perceptual experience of the cup and his background theory of objects and how the world works, then John knows (and may warrantedly assert) A6. Because John is a fallibilist too, his knowledge is not vitiated by his entertaining the epistemic possibility of A6's falsity.

Such is the case for fallible propositions when epistemic agents entertain sceptical alternative possibilities: they are epistemically possible but irrelevant. Why? Because they're destabilizing. One cannot appeal to their improbability, for this is the lesson of the lottery paradox. Therefore, the case is not the same for lottery propositions. Why? Because they're not destabilizing. We do not have grounds to ignore the alternative possibilities where $\neg p$. But, importantly, for both lottery and fallible propositions, the alternative possibilities where $\neg p$

---

48 Hawthorne (2004, pp. 24-28, and elsewhere).

49 See DeRose (1991), for example.

50 One could argue that a strong motivating factor behind Fantl and McGrath (2009) is taking issue with the tension between being a fallibilist but subscribing to EPC.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

are considered as epistemic possibilities. It's just that in the latter, but not the former, such possibilities can be properly ignored and therefore do not vitiate knowledge or assertibility.

NDAT also performs better than salience versions of RATs, such as Hawthorne's, sharing their virtues while avoiding two problems common to such salience accounts. The first of these problems is explaining which alternatives count as salient given the context. Is the (epistemic) possibility that John's cup could have been surreptitiously replaced a salient one? What about whether my car has been stolen from the parking lot? The second problem is the implausible implication that cautious epistemic agents will *ipso facto* know less than incautious agents. The epistemologist may habitually entertain a larger array of alternative possibilities than an incautious layperson. Or an irritating interlocutor may cause an agent to lack knowledge merely by increasing the range of seemingly salient options by mentioning such possibilities as brains-in-vats and quantum-shift coffee cups. But NDAT avoids this problem. Whether an alternative possibility counts as salient on some reckoning is irrelevant to whether the alternative can be rationally ignored or must be eliminated. Rather, what matters is whether the supposition of the alternative would be destabilizing.

7. Implications for Norms of Assertion

In this paper, I have attempted to establish an asymmetry between the epistemic properties of lottery propositions and fallible propositions. One can know fallible propositions such as that I am writing at a computer, but cannot know lottery propositions such as that my ticket will not win in tomorrow's lottery. In order to know some proposition, one must ignore irrelevant alternative possibilities and eliminate, or be in a position to eliminate, the relevant alternative possibilities. Although specifying precisely what constitutes grounds for ignoring irrelevant alternative possibilities has been notoriously difficult, I have argued that an appealing account of the asymmetry can be found in a relative alternatives approach that builds on the differences between destabilizing and non-destabilizing observations, which I have called NDAT.

If the supposition of an alternative possibility would trigger a significant restructuring of our background beliefs about how the world works, then such an event is destabilizing. Discovering that, contrary to every sustained indication, I was not writing at a computer would surely destabilize much of my understanding of the world and the reliability and meaning of my experiences across the board. Discovering that my ticket had won the lottery would be a great surprise, but would require no restructuring of my thoughts on how the world works. Indeed, in doing so I would commit a statistical fallacy. Destabilizing alternative possibilities, though possible, may be properly ignored in the first instance. Non-destabilizing alternative possibilities (such as one's actually winning a lottery) cannot be ignored. In the case of fallible propositions, but not lottery propositions, the alternative possibilities which the sceptic uses as a wedge may be ignored because they're destabilizing. Provided that one also

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

eliminates the relevant (*viz.*, non-destabilizing) alternative possibilities, then one can know such propositions.

In this way, NDAT retains Williamson's basic intuition INT1, explaining why an agent fails to know in the lottery case. Furthermore, unlike Hawthorne's solution, NDAT is able to retain Williamson's second basic intuition INT2. It explains how one can know in the ordinary case. However, NDAT does not depend on identifying the pathology of asserting lottery propositions (such as A1) with the 'merely probabilistic' nature of one's evidence—a move that has problematic implications for Williamson's own account. By rejecting the insufficiency principle, the NDAT version of a Relevant Alternatives approach can remain fallibilist and yet reject knowledge in the lottery case, while allowing for knowledge in ordinary cases.

Although we can now effect a separation between lottery and fallible propositions, such that the latter but not the former are knowable, this is not yet to comment on the implications for norms of assertion. The lottery paradox is one of the principal pillars in the case for KNA: what best explains the unassertibility of (A) and statements like it is that one cannot know that a given ticket will lose. One of my aims with this paper has been to knock down one of the central pillars of support for KNA. As I've argued, the important feature of lottery propositions is that they remain unknowable and unassertible even though they may be true. I've argued that since one can truly believe a lottery proposition, absent of Gettier conditions, the only explanation for the unknowability and unassertibility is that one lacks sufficient (epistemic) *justification* for the proposition. The upshot is that one need not invoke a knowledge norm of assertion to explain the unassertibility: any norm which includes a justification requirement (or what I'll simply call *reasons*) will be able to explain the phenomena. The fact of the matter is that one *never* has the supportive reasons for lottery propositions in order to warrantedly assert them.[51]

If what seems to explain the unknowability and unassertibility of lottery propositions is that agents lack sufficient (epistemic) justification, then invoking a knowledge norm is overkill: all we need is a reasons-based norm, something such as Jennifer Lackey's Reasonable-to-Believe norm (RTBNA) or my

[51] I don't mean to say that an agent lacks all epistemic justification for her belief that her ticket will lose. I merely mean to say that she wouldn't have adequate justification for knowledge (assuming an absence of Gettier conditions) or warranted assertibility. However, in McKinnon (2013b), I've put forward a view on the norms of assertion that may suggest that there are, at least, some contexts where such assertions may be warranted. I remain neutral on the question of whether such an agent would have adequate justification to believe that her ticket will lose. One might think, for example, that she might be justified in believing (at least subjectively), but not warranted in asserting that her ticket will lose. I thank an anonymous referee for raising this issue.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

Supportive Reasons Norm (SRNA).[52] Both a nut-cracker and a sledge-hammer will successfully crack a walnut, but we only need the former. Moreover, the former is best-suited for the task. I think the same is true for something like RTBNA or SRNA *viz*. the lottery paradox data. KNA's explanation depends on the lack of adequate epistemic reasons for warrantedly asserting (or knowing), so a lack of reasons is at the heart of the impropriety of asserting a lottery proposition. Therefore, a reasons-based norm seems able to explain the lottery paradox data at least as well as KNA *and* enjoys a theoretical advantage in rendering the KNA account superfluous.[53]

References

Austin, J. L.
1946 Other Minds. *Proceedings of the Aristotelian Society*, Supp. 20, 148–187.

DeRose, Keith
1991 Epistemic Possibilities. *Philosophical Review*, 100 (4), 581–605.

DeRose, Keith
1996 Relevant Alternatives and the Context of Knowledge Attributions. *Philosophy and Phenomenological Research*, 56 (1), 193–197.

Dodd, Dylan
2007 Why Williamson Should be a Sceptic. *Philosophical Quarterly*, 57 (229), 635–649.

Dougherty, Trent and Patrick Rysiew
2009 Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions. *Philosophy and Phenomenological Research*, 78 (1), 123–132.

Douven, Igor
2006 Assertion, Knowledge, and Rational Credibility. *Philosophical Review*, 115 (4), 449–485.

Dretske, Fred
1971 Conclusive Reasons. *Australasian Journal of Philosophy*, 49, 1–22.

Fantl, Jeremy and Matthew McGrath
2009 *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Harman, Gilbert
1973 *Thought*. Princeton: Princeton University Press.

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

Hendricks, Vincent
2006 *Mainstream and Formal Epistemology*. Cambridge: Cambridge University Press.

Hill, Christopher and Joshua Schechter
2007 Hawthorne's Lottery Puzzle and the Nature of Belief. *Philosophical Issues*, 17, 102–122.

---

[52] Lackey (2007), McKinnon (2013a).

[53] My thanks to Tim Kenyon for his help at various stages of this project.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press

Kindred, Dave
1999 Against All Odds–Hole-in-One Records. http://findarticles.com/p/articles/mi_m0HFI/is_1_50/ai_54390606/ Last accessed December 2010.
Kvanvig, Jonathan
2009 Assertion, Knowledge, and Lotteries. P. Greenough and D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 140–160). Oxford: Oxford University Press.
Kyburg, Henry
1961 *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.
Lackey, Jennifer
2007 Norms of Assertion. *Nous*, 41 (4), 594–626.
Lewis, David
1996 Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
Mayr, Ernst
1991 *One Long Argument: Charles Darwin and the Genesis of Modern Evolutionary Thought*. Cambridge, MA: Harvard University Press.
McKinnon, Rachel
2013a The Supportive Reasons Norm of Assertion. *American Philosophical Quarterly*, 50 (2), 121–135.
McKinnon, Rachel
2013b Getting Luck Properly Under Control. *Metaphilosophy*, 44 (4), 496–511.
Schaffer, Jonathan
2008 Knowledge in the Image of Assertion. *Philosophical Issues*, 18, 1–19.
Smith, Michael
2010 What Else Justification Could Be. *Nous*, 44 (1), 10–31.
Stine, Gail
1976 Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies*, 29, 249–261.
Vogel, Jonathan
1990 Are There Counterexamples to the Closure Principle? M. Roth and G. Ross (Eds.), *Doubting: Contemporary Perspectives in Skepticism*. Dordrecht: Kluwer Academic Publishers.
Vogel, Jonathan
1999 The New Relevant Alternatives Theory. *Philosophical Perspectives*, 13 (13), 155–180.
Whewell, William
1840 *The Philosophy of Inductive Sciences*. London: J. and J. J. Deighton.
Williamson, Timothy
2000 *Knowledge and Its Limits*. Oxford: Oxford University Press.
Wilson, E. O.
1998 *Consilience: The Unity of Knowledge*. New York: A. A. Knopf Inc.

https://doi.org/10.1017/S0012217313000656 Published online by Cambridge University Press