

forthcoming in *The Routledge Handbook of Epistemic Contextualism*  
edited by Jonathan Jenkins Ichikawa

# Lotteries and Prefaces

Matthew A. Benton

University of Notre Dame

## 1 Introduction

Consider a fair coin, which is now flipped. Should you believe that it won't land heads? Presumably not. Consider then a fair die, which is now rolled. Should you believe that it won't come up 6? Or consider two fair six-sided dice, which are now rolled. Should you believe, what is about about 0.97 likely, that they won't both come up 1? Or consider four such dice, now rolled. Should you believe, what is about 0.999 likely, that they won't each come up 1?

Now consider a fair lottery of one thousand tickets, with one ticket guaranteed to win. Should you believe of any particular ticket (what is 0.999 likely) that it will not be drawn the winner? Many people will be tempted to answer this question similarly to how they answered the previous question, about four fair dice each coming up 1: namely, "yes" to both. This is because high probability is widely thought of as a good guide to what one may (and even should) believe. Implicit in this thought is the idea that the epistemically normative notions of what it is permissible (and obligatory) to believe encode some sort of threshold concerning how probable a proposition must be in order for one (acceptably) to believe it.

But if one may acceptably believe that any given ticket will lose, then presumably one may acceptably believe, on identical statistical grounds, that *each* of the lottery's thousand tickets will lose. Suppose one does believe this of each ticket, that it will lose; yet in a typical lottery, one knows

that some or other ticket will win, and so one would know that one of these lottery beliefs is false. This can seem paradoxical in that one seems to have epistemically acceptable grounds for believing each of a conjunction of propositions, even though one knows that the entire conjunction cannot be true.<sup>1</sup>

Similar considerations plausibly arise in preface scenarios. Suppose you have painstakingly researched for writing a long non-fiction book, and have made many claims throughout the book manuscript. Upon finishing it, you write the preface, wherein it is customary to acknowledge with due humility that among the many hundreds of claims you have made, you are bound to have made a mistake: it is highly likely that some of your claims are false (see Makinson 1965). Nevertheless, given that you don't now know that any of your claims are false, or which ones they are if some are false, you stand by each of your claims. And if you stand by each of your claims, it seems that you also affirm their conjunction. This can seem paradoxical in that you the author seem to be claiming, on the one hand, a long conjunction, yet in the same text you claim that it's likely that the entire conjunction is not true.

This article provides a critical overview of how some contextualists in epistemology have approached lottery and preface scenarios. Though lotteries and prefaces raise several interesting issues concerning what one may or ought to believe, I shall focus my attention on what contextualists have said about knowledge and knowledge ascriptions with respect to the preface and lottery scenarios. For on the one hand, it is widely thought that an individual in a preface scenario can appropriately believe (and sometimes know) any particular claim made; and it is widely thought that in a lottery scenario, one can perhaps appropriately believe that a ticket will lose even though the standard judgment is that one does not know it will. Yet on the other hand, when reflecting on the conjunction of such claims, it is also widely thought that cannot know, and perhaps even should not believe, all of the conjoined claims at once.

Evaluating these paradoxes is of particular interest because what we commonly take ourselves to know can seem to depend on the truth of

---

<sup>1</sup>See Kyburg 1961: 197 and 1970 for the original lottery paradox.

lottery outcomes; and our everyday knowledge can also seem to implicate us in being in a position to know, by deductive closure, the truth of lottery outcomes (Hawthorne 2004: Ch. 1). The lottery paradox arguably relies on a principle of deductive closure according to which knowledge may transmit from known premises to a known conclusion, for example:

(MPC) If one believes a conclusion by competent deduction from some premises one knows, one knows the conclusion.<sup>2</sup>

MPC, or something similar, is initially quite plausible.<sup>3</sup> If I know that I won't be able to afford an African safari next year, it might seem that I can know what this logically entails, namely that I won't win a lottery or inherit a huge fortune; or if I know I will be in California in December, it might seem that I can know (what this entails) that I won't suffer an unexpected but fatal heart attack in the interim. But in a wide variety of mundane situations, we are willing to ascribe knowledge of the former claim, yet deny that one could (on similar grounds) know the entailed claim. Contextualists in epistemology offer an initially compelling explanation of this set of judgments, including those involving lotteries, while upholding a principle of closure.

## 2 Contextualism about Knowledge and “Know”

Contextualists about knowledge ascriptions are broadly committed to a theory on which the truth-conditions of “know(s) that” is sensitive to the context in which it is used. The semantic contribution of “know(s) that” to a sentence is determined in part by the speaker (and interlocutors) at a conversational context. Just as indexical terms like “I” and “here” depend

---

<sup>2</sup>Williamson 2000: 117 and 2009. Hawthorne's version of Multi-Premise Closure reads thus: Necessarily, if S knows  $p_1, \dots, p_n$ , competently deduces  $q$ , and thereby comes to believe that  $q$ , while retaining knowledge  $p_1, \dots, p_n$  throughout, then S knows that  $q$  (Hawthorne 2004, 33).

<sup>3</sup>Though sympathetic to closure, I don't intend to discuss the debate here. See Hawthorne 2004: 31–50 and Hawthorne 2014; Dretske 2014; and Williamson 2009. For contextualism and closure, see also Maria Lasonen-Aarnio's contribution to this volume.

for their semantic contribution on the speaker or location of use, or gradable adjectives like “flat” or “tall” depend for their meaning on a standard being assumed by the speaker for what counts as flat or tall given the interests of the conversational context, so the semantic value of “know(s) that,” according to contextualists, depends in part on the speaker’s context of use. One attraction of contextualism is that it can explain the allure of skeptical arguments: they gain traction by subtly shifting the standard for the truth of knowledge ascriptions away from the manageable everyday standard to a much higher standard which is nearly impossible to meet. Thus in high standards skeptical situations, sentences of the form “S knows that *p*” are typically false.

David Lewis’s contextualist approach purports to define knowledge in terms of both the subject’s evidence and which possibilities the ascriber of knowledge is properly ignoring:

(LC) S knows proposition P iff S’s evidence eliminates every possibility in which not-P—Psst!—except for those possibilities that we are properly ignoring. (Lewis 1996: 554; repr. 1999: 425)

Lewis’s definition owes much to the “relevant alternatives” approach to theorizing about knowledge, which must provide an account of what and why certain alternatives count as “relevant.”<sup>4</sup> As such, Lewis’s definition is supplemented by several rules clarifying which possibilities may and may not be “properly” ignored (1996: 554ff.). I shall highlight only four. The *Rule of Actuality* says that “The possibility that actually obtains is never properly ignored; actuality is always a relevant alternative; nothing false may be properly presupposed.” The *Rule of Belief* states that a possibility believed by the subject to obtain is not properly ignored, whether or not she is right so to believe it. The *Rule of Resemblance* maintains that for two possibilities that saliently resemble each other, if one of them may not be ignored, neither may the other be properly ignored. And the *Rule of Attention* states that “a possibility not ignored at all is *ipso facto* not properly

---

<sup>4</sup>See especially Goldman 1976 and Stine 1976; for discussion of such a view with lotteries in view, see McKinnon 2013.

ignored”: those possibilities currently being attended to by the speaker and hearer are not properly ignored.

As written, LC provides a definition of knowledge itself, such that whether a subject *S* knows a proposition depends in part on the nature of *S*'s evidence but also on features of the ascriber's conversational context: thus its final clause refers (with “we”) to *us*, the ascribers, who might be speaking of whether *S* knows. If we take it straightforwardly as a definition of knowledge, it would seem to create (seemingly) absurd scenarios where a single subject can both know and not know a proposition at the same time. For if in *S*'s own context, *S* is considering her evidence and whether she knows *p*, she might, given Lewis's rules, be properly ignoring all the possibilities not eliminated by her evidence, and so according to the definition, *S* indeed knows that *p*. But in another conversational context, some interlocutors are attending to a possibility  $\neg p$ , which is not eliminated by *S*'s evidence, and thus according to the definition, *S* does *not* know that *p*. As such, taking LC as definition of knowledge results in it defying the law of non-contradiction.<sup>5</sup>

But though Lewis himself construes his task as defining knowledge, and although he conducts his discussion in the object-language of what it takes to know, he acknowledges (in his paper's final paragraph: 1996: 566–567) that LC should really be understood meta-linguistically rather than as a definition of knowledge itself. Stated more properly in meta-linguistic terms, Lewis's LC generates a straightforward semantics for “know(s) that” ascriptions (‘AC’ for *ascriber* contextualism), thus:

---

<sup>5</sup>Another potentially absurd result: Is LC committed to the possibility of *S* knowing some *p* precisely because there is no conversation at all in which speakers are considering whether *p*, or speaking of whether *S* “knows” that *p*? Suppose *S* has no evidence concerning *p*, but our conversation does not broach the matter of either *p* or  $\neg p$  (and let us suppose that *p* is actual, and that *S* does not believe either *p* or  $\neg p$ , etc., so Lewis's other rules do not apply). Lewis says that “what is and what is not being ignored is a feature of the particular conversational context” (1996: 559). In the envisioned case then we are ignoring, and properly so, all the possibilities in which  $\neg p$ ; but then *S*'s evidence trivially eliminates every possibility in which  $\neg p$  except for those we are properly ignoring, because in our context, *every*  $\neg p$  possibility is being properly ignored by us. Since the set of possibilities needed to be eliminated by *S*'s evidence is null, then according to Lewis's definition (given its “iff” clause), *S* knows that *p*, despite *S* not believing *p*, and having no evidence for *p* whatsoever.

(AC) X's ascription of the form "S knows that P" is true iff S's evidence eliminates every possibility in which not-P—Psst!—except for those possibilities that X and X's conversational participants are properly ignoring.

Where the ascriber is the subject, that is, where  $S = X$  such that X is speaking of whether X herself knows, we get the result that whether X can truthfully claim that "I know that  $p$ " depends solely on X's evidence and whether X (and X's conversational partners) can properly ignore certain possibilities not eliminated by X's evidence. X can speak truly in self-ascribing such "knowledge" even while another ascriber (say, Joe) can simultaneously, and truthfully, claim that "X does not know that  $p$ ," perhaps because Joe himself is attending to possibilities uneliminated by X's evidence. The semantics given by AC thus offers the truth-conditions for ascriptions and denials of "knowledge" relative to the alternatives being attended to or properly ignored in the ascriber's context; the AC contextualist need not be committed to there being a fact of the matter, independent of such conversational contexts, about whether a subject "knows."<sup>6</sup>

Two other prominent contextualists offer similar semantic accounts of ascriber contextualism. Cohen's (1988) proposal is that a subject S may be said, truly, to "know" a proposition  $p$  relative to a set of relevant alternatives, where the standards that govern which alternatives are *relevant* is a context-sensitive matter: "the truth-value of an attribution of knowledge is context-sensitive" in similar fashion to the way that indexical terms like "I" and "here" are sensitive to their context of utterance. "As such," Cohen says, "one speaker may attribute knowledge to a subject while another speaker denies knowledge to that same subject, without contradiction" (1988: 97). For Cohen, the shifting standards of what is relevant tracks the standard of epistemic justification required of the subject, by the speaker, in order for the speaker's knowledge ascription to be true.

---

<sup>6</sup>As such the law of noncontradiction plausibly does not apply, for on AC there is no fact of the matter, independent of such ascribing (denying) contexts, about whether someone knows. DeRose (2009: 166ff.), another ascriber contextualist, thinks of "knows" on analogy with gradable adjectives such as "tall": there is plausibly no fact of the matter about who or what counts as tall, apart from ascribers' evaluations at a context.

Similarly, DeRose defends a view on which

the truth-conditions of knowledge-ascribing and knowledge-denying sentences (sentences of the form ‘S knows that p’ and ‘S does not know that p’ and related variants of such sentences) vary in certain ways according to the context in which they are uttered. What so varies is the epistemic standard that S must meet (or, in the case of a denial of knowledge, fail to meet) in order for such a statement to be true. (DeRose 2009: 2–3)

For DeRose, the shifting standards govern the truth-tracking “strength of epistemic position” required of the subject, by the speaker, in order for the speaker’s knowledge ascription to be true. For each of the above contextualist semantics, including AC, the underlying idea is the same: contextualists claim that “know(s)” expresses a different relation between a subject and a proposition depending on the ascriber’s<sup>7</sup> context of use, where the conversational context sets the standard for what possibilities may be properly ignored, or what alternatives count as relevant, or how well the subject must be epistemically justified.<sup>8</sup>

### 3 Lotteries

When it comes to lotteries, the standard intuitive judgment is that a subject who believes that a given ticket will lose on the basis of the probabilities involved does not thereby know that the ticket will lose.<sup>9</sup> Given that

---

<sup>7</sup>Notice that Lewis’s ascriber contextualism may be modified into a subject-centered contextualism (SC), by replacing the final clause’s “we” with referent to S and those in S’s context. But Lewis himself, though noting the possibility of a subject-centered approach like (SC), prefers an ascriber contextualism (1996: 561). For approaches that center on the subject (but are thereby not semantically-contextualist), see Hawthorne 2004: Ch. 4, and Stanley 2005: Chs. 5–6.

<sup>8</sup>Contextualists diverge on the mechanism for shifting the standards. For Lewis, certain possibilities previously properly ignored may become salient by the Rule of Attention; for Cohen, the mechanism is a similar rule of salience (1988: 109); for DeRose, a rule of sensitivity is the mechanism (1995: 37).

<sup>9</sup>This judgment is widespread in (and outside of) philosophy. See DeRose, Cohen, and Lewis above, plus Harman 1973: 161 and 1986: 21; Nelkin 2000; Williamson 2000:

setup, the standard semantic judgment, for at least a wide range of conversational situations, is that a denial along the lines of “S does not know that ticket 1 will lose,” is true.<sup>10</sup>

Lewis claims that a contextualist who endorses a semantics given by (AC) along with Lewis’s rules has an apt explanation of why. Such a denial would be true (when true) because the statistical evidence on which S might believe the ticket a loser makes likely a possibility (that ticket losing) which saliently resembles the possibility of every other ticket losing: the same statistical evidence is available for thinking of each and every ticket in a fair lottery that it will lose. But those possibilities, for each losing ticket, also saliently resemble the possibility that the eventual winning ticket (call it ticket  $w$ ) loses; yet the Rule of Actuality says that possibility is not properly ignored. The Rule of Resemblance says that of two possibilities which saliently resemble each other, either both are properly ignored, or neither is properly ignored. Because the possibility that every other (losing) ticket wins saliently resembles the possibility that ticket  $w$  wins, one may not properly ignore the possibility that any particular ticket wins. As Lewis puts it, “These possibilities are saliently similar to one another: so either every one of them may be properly ignored, or else none may. But one of them may not be properly ignored: the one that actually obtains” (Lewis 1996: 557).

One difficulty with this explanation, however, is that it seems to depend on the lottery having a winner: for it trades on the statistical resemblance between the winning ticket and the other (losing) tickets. But many lotteries do not have a guaranteed winner. In such lotteries, on a given day a ticket can be drawn but it may or may not correspond to a winning ticket holder.<sup>11</sup> For these lotteries, Lewis’s view will predict that

---

246–49; Hawthorne 2002 and 2004; and McKinnon 2013, among many others. See Turri and Friedman 2014 for empirical studies of non-philosophers’ judgments, including the non-knowledge judgment, about lottery cases.

<sup>10</sup>The contextualist, of course, is in a position to allow that some such knowledge-denials are false, for there might be contexts in which the ascriber *is* properly ignoring the possibility that a person wins.

<sup>11</sup>Compare many actual lotteries where a computer generates a winning number by randomly selecting seven numbers (or by drawing balls with numbers on them); most days, the numbers correspond to no winning ticket having those numbers.



one fails to know that a given ticket is a loser only on those days when the drawn ticket has a winner. For all Lewis has said, lotteries whose drawings do not guarantee that the drawn ticket matches a sold winning ticket are lottery draws where, on the days where the drawn ticket produces no winner, we do know of each ticket that it loses. Yet the intuitive judgment of non-knowledge about losing tickets even extends to cases where there is no winning ticket (DeRose 1996: 571, fn. 6; Hawthorne 2004: 8, 15). To be worthwhile then, a viable contextualist approach must accommodate the commonsense judgements that, (i) ordinary knowledge ascriptions are largely true in mundane situations; (ii) knowledge *denials* are true in high standard (especially skeptical) situations; and (iii) knowledge denials are typically (or at least can be) true when they concern statistical beliefs about lottery propositions, winner or no.

A contextualist strategy might exploit the fact that understanding what a standard lottery is involves grasping that each ticket (or number) is supposed to have an equally good chance of winning; and if every ticket is such that it is as likely to win as every other, then the possibility that any particular ticket wins should be salient to someone considering lotteries, including any ascriber (denier) considering whether a subject knows that a ticket wins. Given the standard statistical grounds for thinking a ticket will lose, the scenario of that ticket winning remains as likely as any other scenario, and thus it represents a salient possibility<sup>12</sup> where despite one's grounds, one's belief is mistaken. It is not the mere reliance on high probability that is the culprit, but reliance on probabilities structured of a standard lottery which are structured so as to make each ticket equally likely, or at least not appreciably differentially likely,<sup>13</sup> to win. As such, the lottery situation automatically installs an epistemic standard which is not met by relying on the lottery's statistical grounds, and thus knowledge denials of the form "S doesn't know that ticket *x* will lose," which implicitly regard S as relying only on a lottery's statistical belief-forming

---

<sup>12</sup>For discussion of what might make for such salience, and how it might rob one of knowledge, see Hawthorne 2004: 62ff. and 168ff.

<sup>13</sup>Even if some tickets are more likely to win than others, we are still apt to judge that one cannot know that a ticket in such a lottery will lose. Cf. Vogel 1990: 26 n. 8, Hawthorne 2004: 8–9, 15ff.

grounds, are rendered true.

However, Hawthorne (2002; 2004: 94ff.) argues that any contextualist approach to this puzzle will have trouble given that a variety of our everyday knowledge claims, made without attending to lotteries or statistical considerations, can depend for their truth on statistical facts. For some of what I know about you (say, that you and I will meet for lunch tomorrow) may depend on you not winning the lottery (because if you win, tomorrow you will instead be busy collecting your prize and interviewing with the media).<sup>14</sup> But if my knowledge that we will meet for lunch depends on you not winning the lottery, and I cannot know you will not win, it may seem that I likewise cannot know that we will meet for lunch. (On the other hand, sticking to the claim that you do know we'll lunch may seem to license inferring from this knowledge that you *won't* win the lottery!)

The contextualist will insist that such mundane knowledge ascriptions can remain true even though they in some sense depend on lottery-like statistical facts; indeed, the contextualist is in principle committed to the idea that in some ascriber situations, a subject S can be said to “know” that he won't win the lottery, even if S happens to hold a lottery ticket in his pocket. For only when such statistical facts become salient to the *ascriber* will the conversational situation generate a new meaning for “know(s)” such that denying S “knowledge” is true. And the contextualist will maintain that a closure principle is upheld insofar as one may only infer from what one may, at a given context, be truthfully said to “know” so long as the meaning of “know” stays constant throughout the context of the deduction (inferring from “knowledge” had in a lower standards context to “knowledge” had in a higher standards context will be deemed illegitimate). Cases where it seems illegitimate to deduce the truth of ordinary knowledge ascriptions (“She knows that we will meet for lunch”) to knowledge of an entailed truth (“Hence she knows that she won't win the lottery”) are explained by appeal to a context shift rather than a failure of closure.

This approach can seem unsatisfying, however, for it can appear as though situations where lottery probabilities are salient end up reveal-

---

<sup>14</sup>See Harman 1973: 161 for this example.

ing what was true all along: for if many such everyday knowledge ascriptions depend on lottery-like statistical facts, and we truthfully deny “knowledge” when reminded of this, the skeptic (speaking from a high-standards skeptical context) can insist that you never spoke truthfully of having knowledge even when your conversational situation rightly didn’t consider the relevance of statistical grounds.

The trouble becomes more pronounced when our mundane knowledge ascriptions undergo conjunction introduction. For even in situations where lottery-like odds seem to be far from relevant, the contextualist will allow that an ascriber can truthfully ascribe knowledge to a subject, and that ascriber can presumably conjoin this (given MPC) with other knowledge that the ascriber may claim for the subject.

Consider a case from Hawthorne (2004: 94–95). Suppose that Joe believes, knowing his friend Alfred’s financial situation, that Alfred won’t be able to afford an African safari this year. Joe also believes, on similar grounds, that his friend Bertha likewise won’t have enough money for such a safari trip. Joe has 5,000 such friends, of whom he believes this of them on similar grounds. On contextualism, in a situation<sup>15</sup> where we (the ascribers) are rightly not considering lotteries, a relation (call it ‘K’) holds between Joe and each of these propositions about what his friends can afford, such that this relation makes true any claim we might make in that situation that Joe “knows” each of these truths.

But suppose further that Joe himself has deduced and come to believe that each of these 5,000 friends, since they will not be able to afford a safari this year, will not win the lottery this year (assuming MPC). Suppose also that each of Joe’s friends presently hold a ticket in a 5,001 ticket lottery, and as it happens, the winning ticket is held by a non-friend. All of this is unbeknownst to us in our conversational situation. Contextualism allows that nevertheless, the K relation can hold between Joe and these propositions about what his friends can afford, and that we, not needing to attend to the lottery applicable to Joe’s friends, can still speak truly by ascribing, in our situation, “knowledge” to Joe of the entire conjunction.<sup>16</sup> “But,”

---

<sup>15</sup>Assume that Joe is not a part of our conversational situation.

<sup>16</sup>Hawthorne 2004: 83 notes that the contextualist would have to ascribe to a metalin-

Hawthorne objects, “it seems crazy to suppose that in *any* context, ‘know’ expresses a relation that holds between an agent and the proposition that none of 5,000 people will win the lottery when the agent has no special insider information and there are only 5,001 lottery ticket-holders” (2004: 95, italics mine).<sup>17</sup> What is worse, if Joe himself has witnessed the distribution of tickets and sees that only one of the lottery tickets is held by a non-friend, and has deduced from what we say he “knows” and has come to believe that the non-friend will win, then K still holds (assuming MPC) between Joe and this proposition, for “the verb ‘know’ in the original context expresses a relation that holds between Joe and the proposition that his non-friend will win!”<sup>18</sup>

## 4 Prefaces

The lottery paradox gets traction by it seeming (probabilistically) reasonable to believe of each ticket that it will lose, which, when conjoined, would lead one to believe of each and every ticket that it will lose, even though one believes that some ticket will win. As such, it may be reasonable to have an inconsistent set of beliefs. But it also seems (intuitively) that one does not know (given just the probabilities) that a particular ticket will lose, and MPC vindicates this judgment by showing how absurd results would follow from knowing (on such grounds) which tickets are losers.

---

guistic closure principle; I ignore such complications here.

<sup>17</sup>Someone might challenge the force of this objection by pointing out that the vignette loses sight of the context of ascription: we would be ascribing “knowledge” to Joe of the entire conjunction in a low standards context where we aren’t attending to lottery considerations. Against this, note that Joe himself has deduced that none of his friends will win the lottery, when it was overwhelmingly likely that one of his friends would do so; the contextualist says that Joe “K”s that none of his friends will win the lottery even though that was overwhelmingly improbable.

<sup>18</sup>Hawthorne 2004: 96–98 sketches a solution available to the contextualist who wants to respect MPC, which deploys a “New Rule of Belief” in the spirit of Lewis. But he notes that such a solution depends on “hyperactive context-shifting” to change the semantic value of “know” to make false the kind of knowledge ascriptions which arise in the above case.

The preface paradox similarly relies on conjunction introduction, but it differs in that each individual belief from the manuscript, taken alone or in sequence, could each (by commonsense standards) be knowledge. For the preface, the inclination to ascribe knowledge is threatened only once one contemplates conjunction introduction by considering all of the claims made in the manuscript and assessing the set as a whole for error. Thus the notion of risk of error in play in the preface is structurally different from that of the lottery, and the threat to MPC is all the more direct (see Williamson 2009: 4ff.).

Nevertheless, MPC retains some credibility here: for it would be bizarre for a speaker, in a given situation, to go through each claim in the book and ascribe knowledge to the author but then be unwilling to affirm the statement “The author knows all the claims made in the book.” Situations in which one is indeed comfortable ascribing knowledge for each individual claim will likely be situations where one is similarly comfortable ascribing knowledge of the conjunction (as MPC would predict). But by the same token, being unwilling to ascribe knowledge of the manuscript’s conjunction of claims plausibly correlates with thinking that at least one of its claims isn’t known. Worth noting in this connection is the hesitancy of many actual authors to claim in a preface that they *have*, in fact, made a mistake somewhere; they thereby avoid actual inconsistency by noting that they “likely” or “surely” or “must” have made an error somewhere.<sup>19</sup>

However, we needn’t consider actual prefaces to see the force of the paradox. One might know, because they tell you, that some of your friends will be attending the next philosophy conference; on similar grounds you might learn that 100 of your friends will be there (and indeed, all 100 show up). But though we might agree that you can know of each individual friend that they will attend, few will be inclined to agree that you can (by inferring from each such belief) come to know that all 100 will turn up (Hawthorne 2004: 48–49). This correlates with an uneasiness to claim that “All 100 will attend.” But though we can be nervous about making such

---

<sup>19</sup>Very few authors convey their humility by acknowledging (as Makinson 1965: 205 put it), that “not everything I assert in this book is true.” Makinson’s cited example (from R.L. Wilder) trades on a definite article: “the errors and shortcomings to be found herein are not [the] fault” of those who gave Wilder suggestions and criticisms.

conjunction ascriptions (or ascriptions of knowing the conjunction), this uneasiness may be squared with MPC and the idea that the subject indeed can know the conjunction. For we might be susceptible to misleading evidence about the truth of a conjunction without thereby coming to drop any particular belief in the sequence. If it is possible for someone to know that  $p$  while believing that one does not know  $p$ ,<sup>20</sup> then presumably one can continue to know a conjunction while believing (due to the misleading evidence) that one does not know the conjunction. And if knowledge is necessary but not sufficient for proper assertion, this would explain such uneasiness without impugning MPC.<sup>21</sup>

Such maneuvers aside, it looks as though the preface scenario generates a particularly tricky case for the contextualist about knowledge ascriptions. For an ascriber who goes through a long (or even short) sequence of knowledge ascriptions for some subject  $S$  need not feel any pressure of context-shift; but to capture the judgment that one cannot (normally) truthfully ascribe to  $S$  knowledge of its conjunction would require that merely introducing a conjunction ascription somehow shifts the context so drastically as to make the conjunction ascription false. Such a solution could be offered. But trading as it does on inductive grounds for thinking there must be some error in the sequence, the contextualist would need to provide a principled way of distinguishing between salient inductive grounds for thinking a subject has gone wrong specific to a preface-style case, and the salient inductive grounds we generally have for worrying about errors. Drawing that distinction without admitting skepticism looks to be a difficult task.<sup>22</sup>

---

<sup>20</sup>Williamson 2011 and 2014 has argued that one can know  $p$  even though it is improbable on one's evidence that one knows  $p$ ; in such cases, one might rationally believe that one doesn't know  $p$  yet continue to believe (and plausibly know) that  $p$ .

<sup>21</sup>See especially Hawthorne 2004: 50 for this point. On whether knowledge is sufficient for assertion, see Lackey 2011, Benton 2016, and Lackey 2016.

<sup>22</sup>Many thanks to Max Baker-Hytch, Nathan Cockerham, Jonathan Jenkins Ichikawa, and John Hawthorne for comments or discussion.

## References

- Benton, Matthew A. 2016. "Expert Opinion and Second-Hand Knowledge." *Philosophy and Phenomenological Research* 92: 492–508.
- Cohen, Stewart. 1988. "How to Be a Fallibilist." *Philosophical Perspectives* 2: 91–123.
- DeRose, Keith. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104: 1–52.
- . 1996. "Knowledge, Assertion, and Lotteries." *Australasian Journal of Philosophy* 74: 568–580.
- . 2009. *The Case for Contextualism*. Oxford: Clarendon Press.
- Dretske, Fred. 2014. "The Case against Closure." In Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 27–39. Malden: Wiley Blackwell, 2nd edition.
- Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771–791.
- Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press.
- . 1986. *Change in View: Principles of Reasoning*. Cambridge: MIT Press.
- Hawthorne, John. 2002. "Lewis, the Lottery, and the Preface." *Analysis* 62: 242–251.
- . 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- . 2014. "The Case for Closure." In Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 40–55. Malden: Wiley Blackwell, 2nd edition.
- Kyburg, Henry. 1961. *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.
- . 1970. "Conjunctivitis." In Marshall Swain (ed.), *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel.

- Lackey, Jennifer. 2011. "Assertion and Isolated Second-Hand Knowledge." In Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*, 251–275. Oxford: Oxford University Press.
- . 2016. "Assertion and Expertise." *Philosophy and Phenomenological Research* 92: 509–517.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549–567. Reprinted in Lewis (1999): 418–445.
- . 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Makinson, D. C. 1965. "The Paradox of the Preface." *Analysis* 25: 205–207.
- McKinnon, Rachel. 2013. "Lotteries, Knowledge, and Irrelevant Alternatives." *Dialogue* 52: 523–549.
- Nelkin, Dana K. 2000. "The Lottery Paradox, Knowledge, and Rationality." *The Philosophical Review* 109: 373–409.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Stine, Gail C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29: 249–261.
- Turri, John and Friedman, Ori. 2014. "Winners and Losers in the Folk Epistemology of Lotteries." In James Beebe (ed.), *Advances in Experimental Epistemology*. New York: Bloomsbury Academic.
- Vogel, Jonathan. 1990. "Are There Counterexamples to the Closure Principle?" In M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer Academic Publishing.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- . 2009. "Probability and Danger." *The Amherst Lecture in Philosophy* 4: 1–35.
- . 2011. "Improbable Knowing." In Trent Dougherty (ed.), *Evidentialism and its Discontents*, 147–164. Oxford: Oxford University Press.
- . 2014. "Very Improbable Knowing." *Erkenntnis* 79: 971–999.