Jennifer Nagel, University of Toronto
November 23, 2018

The Psychological Dimension of the Lottery Paradox

1. Shifting and controversial intuitions

The lottery paradox exposes some tensions in our natural ways of thinking about probabilities, and in how we think about belief itself. This chapter explores the paradox from a psychological angle, arguing that it arises from the flexibility of our cognitive capacities to represent (and reason about) the empirical realm. A better understanding of these capacities can give us a clearer sense of our theoretical options. Ultimately, I take a broad view of the paradox: in my view, it can be triggered not only by discussion of games with stipulated odds, but by topics of all sorts. However, it will be simplest to start with an example inspired by Kyburg's classic (1961) discussion, in which you hold one ticket in a fair lottery, with odds of (let us say) a million to one, in which the draw has been held but the single winner not yet announced. It is very likely that your ticket has lost, but what is the significance of this high likelihood for the rationality of believing that your ticket has lost? If we insist that a threshold of .999999 is not high enough for rational belief, it may seem we are trapping ourselves in skepticism: surely many of the ordinary things we rationally believe about the world are less certain than logical truths. On the other hand, if we do believe that this ticket has lost, by symmetry we should say the same for any of the other tickets in the lottery, and as long as conjunction of rational beliefs is a rational operation, it seems we would be rational to deduce that all the tickets have lost, in contradiction to our other beliefs about this fair lottery.

As a rational person in this situation, do you believe that your ticket has lost? Do you instead restrict yourself to believing only that it is very likely to have lost? Do you even sense a clear contrast between believing a proposition, and believing that this proposition is very likely? These are genuine questions: we will see that there is a strange controversy about what is intuitive here. Whatever one's own intuitive inclinations might be—whether to believe that the ticket has lost, or to believe only the weaker proposition, or to shift between these stances in different contexts—one may wonder about what explains these inclinations, why others might feel differently, and then also how it would be most rational for us to think.

One of the many notable features of the lottery paradox is that experts disagree not only in their final theoretical analyses of the paradox, but also in their initial intuitive reactions to it, and their sense of how broadly these intuitions are shared. Is it intuitively justified or rational to believe that one's ticket in a large fair lottery has lost, given its high chances of losing? According to some surveys of the literature, *prima facie* intuition is widely taken to speak against outright belief. Julia Staffel summarizes the literature as follows, using 'lottery proposition' to denote a proposition such as 'my ticket in this fair million-to-one lottery will lose':

> The following two constraints on what constitutes rational outright beliefs have been proposed in the literature. They are meant to capture our intuitions about the ordinary notion of belief. (1) It is irrational to have outright beliefs in lottery propositions; (2) It is irrational to hold outright beliefs based on purely statistical evidence. (Staffel 2016, 1725)

Staffel goes on to explain the reasoning behind rejecting the rationality of outright belief in lottery propositions, notwithstanding their high probability. According to Staffel, "This response seems quite intuitive: if I can rationally believe that my ticket will lose, why buy it in the first place? Why not throw it away? Moreover, I can rationally think I might win, which is not consistent with believing I will lose" (Staffel 2016, 1725-6).

Not everyone shares Staffel's perception that our "ordinary notion of belief" makes it intuitively irrational to believe a lottery proposition. John Hawthorne, Daniel Rothschild and Levi Spectre contend that "at least intuitively, it seems reasonable to believe that one's ticket will lose," adding that they take the initial interest of the lottery paradox to depend on precisely this intuition (Hawthorne, Rothschild et al. 2016). In a similar spirit, Philip Ebert, Martin Smith and Ian Durbach look at the position described by Staffel as the standard one in the literature, and characterize it quite differently. As Ebert and colleagues see it, there is widespread acknowledgement in the literature that common sense speaks in favor of belief in lottery propositions: "As well as denying that one can know a lottery proposition based purely on the odds against it winning, some philosophers have gone further and denied that one can even *justifiably believe* a lottery proposition on this basis [they cite Dana Nelkin, Jonathan Sutton, Martin Smith, and Declan Smithies]. Those who endorse this claim, however, rarely describe it as 'intuitive' or widely accepted, and almost universally regard it as something standing in need of substantial argument" (Ebert, Smith et al. 2018, 2). The contrast with Staffel's presentation of our pre-theoretical starting point invites curiosity about what really is intuitive or widely accepted.

There are a number of philosophers in each of these camps, including for example Dylan Dodd (2017) on Staffel's side, and others, back to Henry Kyburg (1961), on Hawthorne's. Indeed, the starting point of Kyburg's argument for the rationality of accepting lottery propositions is that, despite their small chance of being false, "everyone simply accepts" statistical hypotheses of just the sort that Staffel has suggested would be outlawed by "our ordinary notion of belief." In his defense of the reasonableness of accepting that the lottery ticket has lost, Kyburg writes: "If it is objected that this is not reasonable, on the grounds that there is a finite probability that ticket j will, after all, win the lottery, we can answer by pointing out that according to the same line of reasoning, there is a finite probability that *any* statistical hypothesis of the sort that everyone simply accepts, is false" (Kyburg 1961, 197). According to Kyburg, there are many propositions which are highly probable but not certain—having, in his view, "nothing but their probability to recommend them" (Kyburg 1961, 87) —where these propositions are nonetheless simply accepted, or believed outright, by all rational people. His core example of this type of proposition is that the sun will rise tomorrow; surely, he urges, it is an unreasonable form of skepticism to resist outright belief in such propositions.

We can agree with Kyburg's claim that only a skeptic would resist the notion that the sun will rise tomorrow, even while wondering whether this notion is relevantly similar to the claim that a given lottery ticket will lose. It is not obvious how we could pull apart the truths about sunrises that we want to embrace from the truths about lotteries that will get us into trouble: it seems there is no straightforward way to carve out the problematic territory in terms of its logical or probabilistic features (Douven and Williamson 2006). But whatever side we are taking, we can also grant the last point made by Ebert and colleagues: substantial argument would be needed to defend the rational legitimacy of rejecting outright belief in a lottery proposition, whether or not rejection was intuitive. *Any* attempt to resolve this paradox will require substantial argument. But there is something striking about the mismatch in initial perceptions of what seems at the outset to be intuitively rational (or reasonable, or justified—for present purposes I take these notions to be equivalent). The divergence here marks a difference between this paradox and the knowledge version of the lottery paradox, where almost all find it intuitive that we lack knowledge that the ticket has lost, despite strong odds. Somehow, the belief version of the paradox has an extra layer of difficulty about its starting point, even before we search for possible solutions.

One might imagine that the best way forward on the question of what is intuitive would be to turn to formal data collection, and run surveys (on laypeople or experts, or both) on whether

they feel it is rational to believe lottery propositions. Indeed, Ebert and colleagues (2018) have done just this, finding majority support for answers in line with their view of what is intuitive, although interestingly they find a significantly weaker majority for the positive attribution of rational belief than for the negative knowledge attribution. John Turri and Ori Friedman report similar findings, with roughly 80% of their participants taking a person considering purchasing a lottery ticket to be justified in believing that it will lose, while 90% of their participants took this agent not to know the ticket will lose (Turri and Friedman 2014). However, it might be premature to take the question as settled by the majority response to these studies, not least because it is possible that naïve experimental participants may have understood the justification question as asking about the pragmatic (as opposed to the epistemic) rationality of believing that the ticket will lose.

Another possibility is that both sides in the epistemological debate are in a sense right: perhaps there are some ways of speaking or thinking (about lotteries, or about belief itself) that make it seem intuitive to say or judge that a rational person will believe that the ticket will lose, and there are other ways of speaking or thinking that make this seem irrational. A particular way of setting up the question or running a survey may tend to cue one of these ways of speaking or thinking at the expense of another. If this is right, then the question of what is intuitive—either among experts or among laypeople—remains an empirical question, but a question with a more complex answer than we might at first have supposed. The main task in what follows will be to explain and defend the idea that we have diverse ways of speaking and thinking here, giving a descriptive account of their core descriptive features before turning to the normative question of their rational merits. With a better understanding of our resources for detecting and speaking about belief, and perhaps also a deeper understanding of belief itself, we will be better positioned to find a safe path between the threat of skepticism and the threat of odd restrictions on logical reasoning.

2. Speaking of believing, and thinking

If we want to understand our intuitive ways of thinking about belief, one possible source of data could be our spontaneous, everyday use of the verb 'believe', and similar verbs. Some researchers are inclined to make much of this data. Kevin Dorst, for example, argues that linguistic evidence decisively favors the Lockean view, according to which there is nothing more to belief than high

confidence (Dorst 2018). Following Hawthorne, Rothschild and Spectre (HRS), Dorst takes the proposition-embedding sense of the verb 'think' to be the most natural expression of belief in English; Dorst also agrees with HRS that this sense of 'think' is synonymous with 'believe,' on the strength of phenomena such as the incongruity of saying things like, "He thinks that *p* but it's not as if he believes that *p.*" In Dorst's view, although the relevant senses of 'think' and 'believe' express the same mental state, our use of 'think' is an even better guide to our ordinary intuitive under-standing of belief than our use of 'believe', both because 'think' is "more colloquial" and because philosophers use 'believe' in perhaps technical expressions such as 'believe outright' and 'fully believe', possibly contaminating our intuitions about belief-talk with theory (Dorst 2018, 3). When we look at how 'think' is ordinarily used, Dorst contends, we will see that it is utterly routine to equate high confidence with belief. For example, he observes that having heard a forecast that there is a 70% chance of rain makes it seem fine to answer the question of whether it will rain with, "I think it will." Equally, one can say of a lottery ticket holder, "I think her ticket will lose," without any obvious irrationality. In his view, "We are constantly having to figure out what to think based on merely statistical, inconclusive evidence—and outside the philosophy room we do so without pause or compunction" (Dorst 2018, 5). If we apply the T-schema to ordinary language, with a charitable presumption of common rationality, Dorst concludes, we will see that there is nothing irrational about what he takes to be the common practice of believing lottery propositions.

Dorst is certainly right that 'think' is more common than 'believe': in a balanced corpus of written and spoken English, these are the 12th and 50th- most common verbs, respectively (Davies and Gardner 2010, 317), and in spoken language, 'think' is more than six times as common as 'believe' (Davies 2008-, accessed November 13, 2018). Focusing on the first-person, one recent study found 'I think' to be used roughly once every three minutes in unplanned American English speech, and every six minutes in planned speech; it is taken to be the most frequent 'epistemic phrase', while 'I believe' does not figure in the top ten (Kärkkäinen 2010). Given its relative frequency, 'think' has a good claim to being the default nonfactive epistemic mental state verb in English, and, if we restrict ourselves to the sense that embeds a propositional complement, its interchangeability with 'believe' is largely regarded as unproblematic in the mental state attribution literature.[1]

---

[1] It is an interesting question why philosophical discussions focus on 'believe'; I suspect that this choice may be driven by the easier nominalization and by a desire to keep clear of any ambiguities arising from the sense of 'think' that means 'ruminate'.

What is more controversial is whether the nature of belief can be read directly off our everyday use of these verbs. One worry about Dorst's first-person examples in particular is that both 'I think' and 'I believe' are most frequently used not to self-attribute a mental state, but to modulate an assertion. For example, the conspiracy theorist who says, "I think it was Bush who planned the 9/11 attacks" is ordinarily heard as having said something false about the attacks, not something true about his or her own mental state. Note that the 'I think' here could easily be moved to medial or final position in the sentence, as befits an adverbial expression. While it is true that initial 'I believe' and 'I think' can have main clause status, this status is rare, comprising less than five percent of the instances in a recent study of spoken usage (Kaltenböck 2009). Both expressions much more commonly appear as parentheticals with an adverbial use modifying the speaker's commitment, or a discourse-marking use, said to carry little of the phrase's original semantic meaning. There are a number of rival theories about the proper taxonomy here. Gunther Kaltenböck (2010) identifies four different functions of parenthetical 'I think': (1) shield, (2) approximator, (3) structural, and (4) booster. The first of these uses is by far the most common (over 80% of the instances in his 2010 study of parenthetical 'I think'). As a shield, 'I think' has the function of a hedge, or "an epistemic marker expressing the speaker's tentativeness, lack of certainty or commitment with regard to the truth value of its host" (Kaltenböck 2010, 248). As an approximator, 'I think' takes phrasal rather than clausal scope, inserted before precise expressions such as dates and measurements—as in, "the story was in I think the October issue"—and serves "to trigger the need for greater fuzziness or vagueness" (2010, 254). The 'structural' or discourse-marking function of 'I think' is harder to pin down, not least because its scope is unclear. Kaltenböck observes that 'I think' disproportionately occurs amidst disfluencies, surrounded by pauses, backtracking and hesitation sounds (uh, um); in his view, 'I think' often functions as some type of filler, bleached of literal meaning, perhaps signaling the speaker's effort to maintain their conversational turn, or it may serve some linking function between complex topics and comments. The booster function is intriguing: with the right pitch movement and emphasis on the personal pronoun, 'I think' can express heightened rather than diminished commitment to the host proposition. In Kaltenböck's view, the booster function is possible for subjective (for example, aesthetic or political) evaluations ("I think Opera North is always inventive"), where the hedging shield function is invoked for 'verifiable' matters ("I think John is in London") (2010, 264).

Nicole Dehé and Anne Wichmann argue that these function types are distinguished from each other by prosody, both for 'think' and for 'believe'. In their corpus study of sentence-initial *I think (that)* and *I believe (that)*, they argue that initial 'I think' (or 'I believe') has main clause status (they

locate the boost function here) when the speaker stresses the pronoun; it functions as a hedging parenthetical or comment clause when the verb is stressed (sometimes the pronoun is even dropped in this usage), and as a discourse marker when both pronoun and verb are unstressed (Dehé and Wichmann 2010). This approach suggests one possible step towards an explanation for our philosophical divergence over what is initially intuitive. With a difference of emphasis of the sort that is not usually marked in writing, the following sentences can be read as making quite different claims:

(1) I ***think*** this lottery ticket will lose.
(2) ***I*** think this lottery ticket will lose.

In a context where the long odds are public, tentative-sounding (1) sounds good. But because the 'I think' here functions as a parenthetical hedge indicating that what follows is only probable, rather than as a main clause self-attribution of a mental state, there is limited value in the fact that (1) sounds like a truth spoken by a rational person. We cannot simply follow Dorst's advice to reason "according to the T-schema" (2018, 3), and conclude that a rational person believes the ticket will lose, full stop. By contrast, the bolder declaration (2) does make a mental state claim: here 'I think" has main clause status. However, in contrasting the speaker's epistemic position with the hearer's, (2) seems more questionable if the speaker does not have private information warranting a deeper pessimism than the publically known low chance of winning. If Dorst is focusing on cases like (1) while Staffel is focusing on cases like (2), they are responding to different stimuli, although it remains to be seen what the significance is for their underlying disagreement about rational belief itself, and what happens when we switch to the third-person perspective.

There are a number of theories about how and why the phrases 'I think' and 'I believe' have evolved to serve this variety of functions. It is widely noted that there is something oddly redundant about overtly tagging a statement with 'I think': as Graham Ranger puts it, "the operation involved in *I think* is implied in the act of assertion, in *any* case, as part of the enunciative scenario. The question therefore is to determine what the *explicit* presence of the string *I think* adds to an utterance" (Ranger 2018, 282). In Ranger's view, typical contexts of assertion—for example, discussions of whether John is in London—involve a shared default orientation towards knowledge that will set the speaker's overt use of 'think' in contrast with the stronger 'know', generating a tentative effect. Meanwhile, in more subjective evaluative contexts, such as discussions of music or politics, overt '*I* think' is understood in contrast to '*they* think' or '*you* think', boosting the speaker's commitment to what is said (2018, 281-2). This reading fits with Dehé and Wichmann's

observations about the difference between prosodic emphasis on the verb and subject, and Kaltenböck's position on the contexts where 'I think' serves to boost rather than hedge. However, we have at best the start of a theory here, leaving much unexplained, including cases like our (2), a boosted but not especially subjective or aesthetic evaluation. In summation, it remains controversial exactly what range of functions is served by 'I think', but it is widely agreed that these functions are various, and that many of them involve bleaching of the literal mental state meaning of 'think' (on this point, see also Van Bogaert 2011). Where Dorst argued that ordinary usage of 'think' could be a direct guide to the nature of belief, it appears that considerable caution is in order.

3. The difficulty of belief attributions, and beliefs themselves

It might seem that we could make things easy by switching to third-person attributions. Focusing on a third party who knows that it is very likely that the ticket has lost, we can ask whether this person would be rational to believe that the ticket has lost. It is possible we could still suffer some contamination from formulaic uses of 'I think' to mark weakened commitment, for example if in evaluating the mindset of a third party we imagine them saying "I *think* this ticket will lose", or if we imagine ourselves in their shoes, and ask whether we could reasonably say such a thing. Furthermore, there is also an evidential use of third-person mental state language. Mandy Simons (2007) notes that one might answer the question, "Who was Louise with last night?" with something like, "Henry thinks that she was with Bill." To choose this answer rather than the simpler "She was with Bill" is one way the speaker could signal diminished commitment to the embedded clause, Simons observes. Given that it is the embedded clause that is an answer to the question, this clause has main point status: if we are concerned with Louise's company, not Henry's state of mind, our talk of thinking serves chiefly to mark a tentative attitude to the key content. Simons does hold that the speaker is still making an assertion about Henry's mental state here, but it is an open question whether this evidential use of the mental state verb clouds the significance of our intuitions about third-person belief attributions.

But even if we do not suffer this kind of linguistic interference, explaining how we intuitively judge the beliefs of third parties remains a tricky business. There is no question that belief attribution and evaluation forms a major part of mature human social intelligence; indeed, it is argued that the capacity to attribute beliefs (as opposed to just states of knowledge and ignorance) lies at the very core of what sets us apart from other socially intelligent animals

(Tomasello 2018). Spontaneously, and often without any deliberate effort, we see other agents as having beliefs, sometimes misconceptions or accidentally correct opinions, sometimes states of mind we evaluate more positively. We rely on a rich set of cues in belief attribution, including signs of what others can perceive and have perceived in the recent past, what they say, and how they act. This is a complex matter, not least because these cues can point in different directions (for some interesting complications, see Nyarko and Schotter 2002, Schlag, Tremewan et al. 2015).

Because the processes that support spontaneous belief attributions are not generally open to introspection, it takes significant work to develop a theory of how we are attributing and evaluating the beliefs of others. In fact, there are even questions about the nature and extent of our access to our own beliefs. According to one prominent theory of mental state attribution, Peter Carruthers's Interpretive Sensory Access (ISA) model, even in the first-person case we gauge belief indirectly, from behavioral cues (Carruthers 2011). While the range of data we can draw on in the first-person case is different, because it includes inner as well as outer speech, and direct access to phenomenal states such as hunger and pain, the ISA theory maintains that we apply the same mindreading processes to ourselves as to other people in attributing propositional attitudes: for example, querying (inner or outer) speech, and interpreting what we hear or find ourselves saying in our minds. According to the ISA, mindreading involves rapid processes of unconscious inference. We instinctively model ourselves as being self-transparent or having unmediated access to our own propositional attitudes, and these simplified models are generally unproblematic because we are largely accurate in our self-attributions; however, the ISA model maintains that we can be mistaken about what we believe or want. Carruthers supports his theory by reviewing a large body of empirical literature on manipulations and special circumstances that can make people misunderstand their own beliefs and motivations, separating their self-conception from the real underpinnings of their agency (Carruthers 2011, ch. 11). If the ISA model is right, we are able to self-attribute mental states not by direct conscious awareness of the propositional attitudes themselves, but through an interpretive process anchored in our awareness of the effects of these attitudes in (inner) speech and outer behavior. I consider the ISA model to be the most plausible explanation of our existing empirical evidence, including phenomenological evidence, and I also think that it can support a new way of making progress on the lottery paradox, so I will be assuming it in what follows. Those who are initially unconvinced by Carruthers might see some progress on the lottery paradox as a reason to reconsider the merits of the ISA; failing that, there may still be some observations of interest in what follows, as not everything in my approach to the problem will depend directly on the ISA.

Meanwhile, there is a delicate relationship between the nature of belief and the evidence we can draw from the behavior of agents. If someone keeps a lottery ticket, rather than throwing it away, one could read this either as a sign that they do not believe outright that it will lose, or perhaps instead as a sign that their practical reasoning is guided by degrees of belief rather than binary states, and they have rationally calculated a positive expected value in holding onto the ticket (on the significance of this line of thought to the knowledge version of the lottery paradox, see Douven 2008). The latter way of reading the action leaves it open that the rational agent might believe the ticket will lose, assuming high probability is sufficient for outright belief.

Because the relationship between belief, credence and action is so heavily contested, it seems that an easier place to start investigating the attribution and evaluation of binary and graded mental states could be the domain of explicit reasoning, where it is easier to see what we are working with. Does the rational person reason explicitly in terms of binary or graded states? Here again, the answer may not be simple: arguably, we can do both. In a series of recent papers, Jonathan Weisberg has defended the view that both binary and graded representations figure in rational thought, and neither is reducible to the other (Weisberg 2013, Weisberg forthcoming). Depending on the circumstances, we can engage either in thinking which takes categorical claims as input, such as syllogistic reasoning or processes such as Elimination By Aspects, or we can engage in probabilistic calculations, assigning subtle credal weights to relevant factors. Indeed, as Weisberg has pointed out, some interesting general models of rational judgment, such as the Evidence Accumulation Model (Newell and Lee 2011), take both binary and weighted input at different stages of processing.

Explicit reasoning is also a domain in which intuitive evaluations of the rationality of belief can separate from intuitive evaluations of knowledge: the person who engages in valid reasoning from premises she could not have known were false might still be evaluated as rational, even if she fails to attain knowledge from her reasoning. Special attention to this feature of inference should help us understand the relationship between the knowledge version of the lottery paradox and the rational belief version. Our first step will be to look at the emergence of both forms of the paradox in contexts of switching from thinking unreflectively to engaging in explicit reasoning, and here we will turn to an example involving subject matter beyond lotteries.

4. Intuitive and reflective cognition

Who is the current pope? For an easy trivia question like this, an answer comes to mind without conscious effort or reflection.[2] At the time of writing, the answer, which we will call proposition *p,* is that *the current pope is Pope Francis.* Notoriously, however, it is possible for me to induce conscious reflection in the vicinity of proposition *p* in a way that can trigger a different frame of mind about it. For example, we could review mortality statistics for the relevant demographic: according to the World Health Organization's most recent report, men of Pope Francis's advanced age in his world region have a background annual mortality rate of 6.9% (WHO 2016). Other things being equal, one's credence that any such man has not died in the last hour should be about .9999918384. On the question of whether other things are equal, one might also reflect on the risk of assassination and airplane crashes for a very prominent and somewhat controversial leader who travels frequently. It is extremely unlikely that anything untoward has happened to Pope Francis since the last time our news feeds were updated, but he is a man, all men are mortal, and few of us know the precise hour of any remote individual's passing. Having accepted that it is not impossible that Francis has died in the last hour, we might then attend to the further fact that if he has died, we have entered a papal interregnum, making it no longer the case that Francis is the current pope.

With these facts in mind, let us revisit our attitude to the proposition *p*: *the current pope is Pope Francis*, where this is explicitly recognized to entail that it is not the case that this man has died in the last hour. How are we inclined to think about this proposition now? As philosophers at least since Jonathan Vogel (1990) have observed, when we reflect this way, the proposition no longer seems to be known. Before we turn to the question of whether our current frame of mind allows rational belief in the proposition, we can examine the nature of this frame of mind, and the reasons why we tend to feel knowledge has lapsed.

If my example worked as intended, the initial presentation of a question about the identity of the Pope triggered an answer immediately, without conscious reflection on any supporting (or potentially undermining) considerations. The second time the question was posed, it was after attention had been directed to a series of premises presented as relevant to the issue. There is a psychologically significant contrast between the two ways of thinking employed in these successive moments, ways of thinking now commonly labelled as Type 1 (intuitive) and Type 2 (reflective)

---

[2] Any readers who encounter resistance with this example should construct another they find easier; any readers within the pope's immediate ambit should construct an example within the realm of background trivia rather than current experience.

cognition.[3] They differ in a variety of ways, with intuitive processes operating in a swift and autonomous fashion, where reflective thinking is generally slower and often controlled. It is now widely thought that the ultimate basis of the contrast concerns the involvement of our limited capacity central working memory resource (Evans and Stanovich 2013, Carruthers 2015). Working memory keeps contents in mind or consciously available during complex thought, enabling 'cognitive decoupling', or representations of alternatives, or hypothetical states of affairs potentially at odds with what the subject is currently experiencing (Baddeley 2007, Stanovich 2011). To take a classic example (inspired by Sloman 1996), when intuitive cognition solves a slightly shuffled anagram like "auotmonous", the answer ("autonomous") springs to mind with no consciousness of intermediary stages; when reflective cognition solves a more complex and deeply shuffled anagram like "tebardelie", there are successive conscious presentations of possible solutions in which the letters are re-ordered, where each such presentation is evaluated intuitively as a non-word, and further permutations are brought to mind, until the solution is found ("deliberate"). The two types of thinking are deeply integrated: reflective cognition can answer new questions whose answers are not automatic for us, but it does so through organizing a series of cycles of intuitive cognition. Although reflective cognition is sometimes characterized as "conscious", both types of thinking incorporate a substantial amount of processing below the radar of consciousness, such as the processes that govern the allocation of attention and the retrieval of relevant memories (Evans 2009). Some problems can be answered through either way of thinking, but partly because reflective thinking is effortful, our default is to answer intuitively or take for granted the most plausible state of affairs, expending the effort to consider alternative possibilities only when prompted (Evans 2007, Stanovich and West 2008).

The distinction between categorical and probabilistic judgments is orthogonal to the distinction between intuitive and reflective cognition: the suggestion here will not be simply that intuition is always categorical, or that reflection on its own cues probabilistic thinking. Either type of thinking can be engaged with categorical content, for example in solving arithmetical problems either automatically or deliberately. Although very easy problems like single-digit addition are typically solved in intuitive cognition by the automatic production of an overlearned answer, they can be solved through reflective thinking, for example by application of deliberate counting procedures. Either type of thinking can also be engaged to thinking about probabilities, with some differences worth watching. Intuitive thinking is more common with verbal scales of likelihood

---

[3] I am following the recent joint work of Jonathan Evans and Keith Stanovich (2013) in using the adjectives "intuitive" and "reflective" to denote Type 1 and Type 2 cognition, respectively.

('very likely'), and switching from these qualitative scales to numerical scales ('80% chance') tends to shift us towards reflective thought (Windschitl and Wells 1996). However, it remains possible to engage in reflective cognition with qualitative probabilities, for example when the premises of one's explicit syllogistic reasoning include modal terms such as 'likely'. Still, this is not to say that either type of cognition can be used for any problem: reflective cognition becomes mandatory when we need to engage in controlled reasoning about precise numerical representations of likelihoods, given novel problems to which no answer has previously been learned. Differences in our spontaneous epistemic evaluations of intuitive and reflective cognition can play a role in explaining diverse responses to the lottery paradox if subjects can differ in presenting the key proposition in the paradox as generated by one form of cognition or the other.

One final psychological factor that seems relevant to the lottery paradox is the distinction between 'verbatim' and 'gist' representations of probabilities (Reyna 2004). Precise verbatim representations enable detailed calculations, but they are hard to retain and are handled through slow serial processing; fuzzy gist representations may distort fringe cases but are more easily retained and can be handled through swifter parallel processing. In the rougher intuitive format of gist representation, "very small risks are edited to 'essentially nil risk'" (Reyna 2004, 63). Something like gist representation may be behind the phenomenon of insurance resistance for low-probability harms: for example, in simulated management games, people are (surprisingly) more than twice as likely to purchase $500 insurance for a .25 probability of a $1,980 loss than for a .002 probability of a $247,500 loss (Slovic, Fischhoff et al. 1977). If precise and controlled reasoning were in effect here, one would not expect to see a lowered tendency to secure coverage for the more catastrophic prospect of equal expected value. The intuitive tendency to write off tiny probabilities as zero also means that we respond differently to presentations of equivalent risks over different time horizons. In one classic study of seat belt use, participants were presented with either the increased risk of disabling injury for travel without a seat belt for each trip taken (.00001) or the equivalent risk over a lifetime of driving (.33). Only 10% of those exposed to the single trip statistics indicated that their seat belt use would change as a result, versus 39% of those exposed to the lifetime statistics (Slovic, Fischhoff et al. 1978). While these particular examples are generally taken as evidence that gist representations are normatively problematic, Reyna argues that reliance on gist often increases the accuracy of reasoning: where children and novices can be misled by incidental features of verbatim representations, adults and experts in fields such as cardiology develop stronger capacities to represent risks intuitively, typically increasing the consistency of their judgments in doing so (Reyna, Lloyd et al. 2003). This observation is in line

with a larger program concerning the ecological validity of heuristic thinking: given limitations in time, attention and computational skill, there are many circumstances in which human beings reasoning heuristically on the basis of a subset of available information will be more accurate than those attempting complex procedures to weigh all available evidence in full precision (Gigerenzer and Brighton 2009). The moment of making a decision about the purchase of insurance is not the time to discard a remote possibility as the kind of thing that doesn't happen, but in daily life it could be taxing to the point of paralysis to try to make ordinary judgments with maximally precise weights attached to all conceivable possibilities, and overall higher accuracy in judgment could very plausibly be secured through a general practice of reliance on gist representation. Indeed in many situations we have no choice but to rely on gist: verbatim representations (of probabilities, or of situations more broadly) are hard to retain, and in practice as time elapses, gist representations tend to dominate in recall and assertion (Goldsmith, Koriat et al. 2002).

5. The knowledge version of the paradox

We are now in position to return to our question about the pope. In an earlier paper (Nagel 2011), I argued that shifts between intuitive and reflective cognition help to explain the knowledge version of the paradox: someone who initially knows that Francis is the current pope might be unable to come to know that Francis has not died in the last hour, even when she recognizes that this latter proposition is entailed by the former. Focusing on the question of whether Francis might have died in the last hour can undermine one's prior knowledge that he is the current pope, I argued, either by destroying this knowledge through weakening confidence, or simply by thwarting access to it in occurrent judgment, making it impossible for her to use this prior knowledge as the premise of an appropriate deduction. It was a starting point of this earlier paper that either kind of cognition can yield knowledge: paradigm sources of knowledge include both intuitive processes such as face recognition and reflective processes such as sound deduction. In fact, notwithstanding the psychological differences between these modes of thought, what makes a judgment of either type count as knowledge could ultimately be the same across the board (for example, that one is arriving safely at the truth). Even so, it is sometimes possible for a person to know a proposition only when she judges that proposition in one manner rather than the other: for example, if she has unreliable intuitions in a given domain but is capable of accurate calculations when she thinks reflectively, or sound intuitions but unreliable powers of reflection.

Pairs of judgments like those about the pope—this earlier paper labels these as instances of the "Harman-Vogel paradox"—exploit a special kind of disparity in what we can judge intuitively and reflectively. What is special about the propositions that feature in the Harman-Vogel paradox is that the subject is asked to make reflective judgments that question some standing presupposition behind the corresponding intuitive judgment. Ordinary intuitive judgements about generally stable environmental facts like the identity of world leaders are made on the basis of semantic memory. In answering questions about these facts, we typically presuppose that what we recall was originally anchored in knowledgeable testimony, and that the world remains relevantly as it was when we came to know what we now recall. When the appropriate source and world stability conditions are taken by an evaluator to be met, intuitive judgments easily register as instances of knowledge. But these are presuppositions as opposed to conditions consciously available to the subject: as subjects, we are typically unable to recall to consciousness which original source furnished us with any particular trivia fact we recall intuitively, let alone evaluate the quality of that source or collect consciously available materials to support non-question-begging explicitly reasoned judgments about the relevant stability in the environment.

The question introducing our example ("Who is the current pope?") presupposed that we were not in a papal interregnum. When we were invited to consider the answer to that question ("The current pope is Pope Francis") a second time, this presupposition was no longer available, given our active consideration of the hypothetical possibility of Pope Francis's recent demise. Working memory had been charged with the task of contemplating this hypothetical possibility, and we were switched into a reflective mode of thought, in which we were obliged to search for consciously available reasons to deny the problematic focal possibility. Thinking reflectively, we become self-conscious about the timing of what is retrieved from memory, and we can recognize that what is consciously available does not enable the construction of any good explicit argument against a very recent change in the world that would undermine our claim to know our key proposition p (that the current pope is Pope Francis). When we are challenged to produce an explicit argument in support of the judgment about the current pope that we initially made intuitively, we have no choice but to survey the materials consciously available to us, even though the contents of consciousness played no role in producing the original unreflective judgment. We cannot put forward the truth of that judgment as a persuasive argument in its own support: this would be manifestly circular reasoning. It remains possible for a person to judge that p is the case in such circumstances, but not by valid reasoning from premises that are not in question. One could, for example, think to oneself that Pope Francis simply must be alive, because one is very much

hoping to see him in Vatican City next Sunday, but such reasoning has no claim to a positive epistemic evaluation.

One might wonder why exactly we can no longer make (and positively evaluate ourselves as making) our original unreflective judgment, once worries have been introduced. The distinction between reflective and unreflective judgment is a matter of their causation, but we do not have direct access to the etiology of what comes to mind in any mode of thought. When we engage in reflective cognition, we have conscious access to a sequence of contents, and the global broadcast of these contents in consciousness triggers the activation of multiple mental modules that respond to the relevant content, producing further conscious contents such as visual imagery and inner speech. While these contents are conscious, the manner in which we react to them involves multiple unconscious and involuntary processes, including processes that interpret and evaluate inner speech, and processes that intuitively construct and evaluate arguments (for a review of the literature on these processes, see Mercier and Sperber 2017). This process of reasoning can also produce new beliefs, as intuitive modules for reasoning pick up logical patterns among the contents we assert in inner speech, producing conclusions we decide to accept. Arguments presented in inner speech are subject to the same type of spontaneous evaluation that we apply when hearing the voice of another person. We cannot by sheer force of will make something said in inner speech (such as, "Francis *is* the current pope") register as the manifestation of an unreflectively formed belief. Once we have started to reason explicitly on a question, as we will be triggered to do by rehearsal of the precisely represented mortality statistics and the question about the hypothetical possibility of his recent death, we will find ourselves evaluating the things we hear ourselves saying as part or product of our reasoning. If no good explicit reasoning can support a claim which answers the question we are aiming to settle, then as long as we are in a mode of thought that is weighing the epistemic merits of explicit reasoning, it will seem we lack knowledge of the claim in question.[4] If the ISA model of indirect (or non-transparent) access to mental states is right, it is possible that we actually retain knowledge of the key proposition throughout this manipulation, but

---

[4] Although the original Harman-Vogel cases involved a switch from intuitive to reflective cognition, I think the paradox can be extended to any case of switching upwards to a more self-conscious mode of thought, for example, moving within the domain of reflective cognition from plain syllogistic reasoning to reasoning about one's syllogistic reasoning. One ordinarily trusts one's rational capacities unselfconsciously, although these capacities are structured in a way which is largely shielded from introspection: we are typically aware of successive contents but not directly of the nature of the processing that alerts us to inconsistencies or enables the deduction of sound conclusions (on this point, see Evans 2009). But it is possible to construct cases in which we have doubts about those processes, for example on receipt of misleading information to the effect that we have consumed a drug that impairs logical reasoning (as in Christensen 2010). Detailed discussion of these cases is however a project for another occasion.

are unable to register ourselves as knowing when we switch to the self-doubting reflective mode of thought. If our grasp of our own propositional mental states works through indirect processes, such as querying inner speech, our inability to construct a good argument will bar us from producing an apparently knowledgeable occurrent judgment that *p* is the case. We cannot intuitively evaluate the propositional attitude by being directly conscious of the attitude itself, assuming that the ISA is right. However, knowledge itself could well be a state which can endure even when its typical manifestations are temporarily blocked: if so, the intuitive sense of knowledge loss is misleading.

6. The belief version of the problem

The claim that Francis is the current pope (*p*) is a claim that could be made on the basis of either intuitive or reflective thought: I have argued that it will typically come across as knowledge if made intuitively, assuming that the subject learned this fact from a knowledgeable source, and the world conditions have been appropriately stable. In contexts where these presuppositions have been challenged, efforts to reach *p* reflectively by sound and persuasive reasoning will be unsuccessful, and we will sense a lack of knowledge. Whether belief in *p* is in these contexts intuitively rational is a subtler question: the apparent rationality of a belief depends not only on the content of what is believed, but also on the manner in which the subject is arriving at that content. Our initial statement of the problem left it somewhat open how the subject is thinking. After worries have been raised about the possibility of a very recent death, and the mortality statistics have been taken on board, one might imagine a subject reasoning as follows:

> (R1) I have not checked the news in the last hour. Given the pope's age and regional mortality statistics, he has a .0000081616 chance of having died in this past hour.
>
> I have no further information to settle the question of whether Francis is still the current pope, beyond his age and regional mortality statistics.
>
> 1.0 − .0000081616 = .9999918384.
>
> Assuming he was alive an hour ago, the pope has a .9999918384 chance of being alive now. If he is alive, Francis is the current pope.
>
> Therefore, Francis is the current pope.

This way of reasoning looks problematic: in particular, the conclusion that seems best warranted by this effort at precise reasoning is not that Francis is the current pope, but that there is at most a .9999918384 chance that Francis is the current pope.

However, this is not the only way we could imagine our subject as reasoning, on hearing the mortality statistics. We do not have to see the subject as attempting to demonstrate that Pope Francis is alive on the basis of precise reasoning from the statistical information provided. We could instead imagine the subject recognizing that there is a small chance that the pope has very recently died, but then stepping back to a larger perspective:

> (R2) Last time I checked, Francis was the pope. If he had died, other than in the last hour, I would have heard.
>
> Given Francis's age and regional mortality statistics, he has a .0000081616 chance of having died in this past hour.
>
> It is very unlikely that the pope has died in this past hour.
>
> It is generally reasonable to assume that nothing very unlikely has happened.
>
> It is reasonable to assume that the pope is alive, and, assuming he is alive, Francis is the current pope.
>
> Therefore, Francis is the current pope.

This argument is incomplete in a number of ways. For example, there is no mention of the possibility that the regional mortality statistics are inaccurate, or the possibility of Francis's having relinquished the Papacy, and so forth. However, (R2) is not obviously problematic in the way that the precise argument (R1) was. We could even see this argument as a reflection of the thinking behind a deliberate adoption of a gist representation: note the shift from numerical chances to the qualitative "very unlikely". In situations in which the pope has in fact just died, reasoning of this type will obviously not secure the subject knowledge of the conclusion. But even if we consider the subject who reaches a false conclusion by reasoning in this manner just after the pope's demise, we can allow that the subject was relatively blameless, relying at the crucial juncture on an assumption she could not have known to be false. The argument works as a partial defense of a style of thought that will generally produce accurate answers to questions of the type under consideration, and that overall epistemic rationality of thinking in this manner is one way of defending the subject's

epistemic rationality in holding this conclusion. The argument does not serve to defend the legitimacy of relying on gist in all circumstances: for example, in the context of a much larger argument stringing together judgments about many world leaders, the thought that it is generally reasonable to assume nothing unlikely has happened would not so easily warrant the legitimacy of a series of assumptions about each of many hundreds of particular world leaders, where we are actively working to conjoin these assumptions in a pattern of precise reasoning.

By contrast with the case about the pope, one interesting feature of the original lottery case is that the very idea of a ticket in a fair lottery carries with it a suggestion of chances and precise reasoning. One might easily have thoughts of Pope Frances without reflecting on mortality statistics, but thoughts of lottery tickets with specified odds naturally demand verbatim representation and reflective cognition. This feature of lottery tickets plays a key role in Igor Douven's (2012) treatment of the paradox. Douven holds that high probability generally suffices for rational belief, but he argues that it is irrational to believe outright that one's ticket will lose, because such a belief would violate principles of what he calls 'epistemic hygienics'. Just as we can say things that are strictly speaking true but misleading to others, Douven argues, we can also form beliefs that are strictly speaking warranted but potentially misleading to our future selves: rationality requires us to avoid forming beliefs that are risky in this way (Douven 2010). In a standard context where the high odds of loss are publically known, one cannot assert "your lottery ticket will lose" just on the strength of those odds, Douven argues, because to do so implicates at least weakly that you have inside information about the lottery outcome. A cooperative speaker will not chose to make the categorical statement about loss unless he wishes to suggest that he has information going beyond the known odds; if the cooperative speaker simply wanted to remind the hearer of the low chance of winning, Douven observes that he could easily say something weaker, such as "your ticket is even less likely to win than you may think" (2012, 359-60). If we think of belief formation as a type of communication with one's future self, inscribing a particular sentence in one's belief box for later retrieval, then given that one's awareness of the past grounds of a particular belief may fade over time, it is also risky to encode a plain categorical belief to the effect that one's ticket will lose. At the point of recalling that categorical belief, Douven argues, we risk gaining the impression that we must have had inside information about the lottery warranting a higher degree of certainty about the loss than we would have had from our standing awareness of the long odds of lotteries.

Several objections come to mind. First, if one openly maintains that high probability is sufficient for outright belief, then this does something to diminish the sense that believing one's ticket will lose is really stronger than believing one's ticket is very likely to lose: given the perceived weakness of the categorical belief, it is not entirely clear why exactly it implicates the possession of inside information. True, there is some redundancy in having not only beliefs about the probabilistic setup of the lottery but also the further belief that one's ticket will lose, but in a context where forgetting is expected, some redundancy seems fine, and clarity about the weakness of belief should to some extent weaken the impression that one must have had inside information, in a context where one happens to recall just the belief about loss. Second, one might wonder about the picture of belief as inscription of a particular sentence, to be recalled verbatim at a later time: it seems that our capacity for verbatim recall is generally limited to short durations (Goldsmith, Koriat et al. 2002, Britt, Kurby et al. 2007, Reyna, Corbin et al. 2016). Over time, when one has forgotten the grounding of one's original judgment, it is likely that one will also have forgotten its original verbatim formulation, and recall only the gist, which on Douven's theory should not distinguish between loss and very likely loss.

However, there is something in Douven's approach which is very plausible: he notes that the evaluation of what we say or believe seems to depend in part on how it compares to the field of rival propositions we might have chosen. The most natural way in which these fields of rivals can differ is as a function of the various ways in which one might come to a belief that the ticket will lose. For example, a person might be represented as arriving at the judgment that the ticket will lose on the basis of questionable precise reasoning:

(R3) 999,999 out of the 1,000,000 tickets in this lottery will lose.

Therefore, this particular ticket will lose.

As I am evaluating this way of thinking, if the arithmetic of the set-up provokes me to engage in precise reasoning that the ticket has a .999999 chance of losing, then it looks to me like a mistake for the subject to endorse the absolute conclusion that the ticket will lose, over that more precise rival. Framing things like this will move us towards Staffel's intuitive sense of the problem: the rational person does not outright believe that his ticket will lose. To speak of *outright* belief here, is perhaps to underscore the way the subject is represented as decisively asserting the conclusion of this effort at precise reasoning. If Timothy Williamson (2000, ch. 11) is right that knowledge is the norm of assertion, then our intuitive sense that the subject fails to know that his ticket will lose

arguably plays a role in our discomfort with the subject's assertion of this conclusion. If we take the possibility of direct assertion of this kind to be the thing we should check for in belief attribution, and precise reasoning as our model for ideal belief formation, then we are well on our way to the theory according to which knowledge is the norm of belief.

However, there are other ways of reasoning about the lottery. One might, for example, just highlight the contrast between winning and losing:

(R4) 999,999 out of the 1,000,000 tickets in this lottery will lose, and only one will win.

It is vastly more likely that this ticket will lose than that it will win.

It is reasonable to expect that this ticket will lose; so, I believe it will lose.

If the choice is between winning and losing, rather than winning and being .999999 likely to win, the rational person picks loss. This argument shifts into qualitative talk of probabilities in the second line, nudging us towards a gist representation of the outcome. The conclusion does not state baldly that the ticket will lose, and if it did, we might feel more resistance to it, especially if Williamson is right about the norm of assertion. However, there are multiple ways of attributing a belief that *p* to a subject, and hearing the subject declare that *p* is the case is only one of these ways. In this case, the subject moves from a judgment about reasonable expectation to a self-attribution of belief, rather than a frank embrace of the key content. Still, we can read this self-attribution as good evidence for what the subject actually believes, in the sense that it is the kind of declaration that sets up systematic expectations for how the subject will plan and act in the future. This subject is not figuring out what to believe on the basis of precise numerical reasoning—indeed the subject might appear less rational if she were represented as embarking on a mission of precise numerical reasoning here—but human courses of action are not always driven by precise numerical reasoning. For many of the purposes behind the formation of beliefs, the subject's way of thinking here is entirely sensible. If we represent subjects as thinking this way about the lottery, then we can shift towards Hawthorne's sense of what is intuitive.

The person who believes that her lottery ticket will lose may be reasoning precisely but poorly, or in a qualitative, gist-based manner, and well. If one has an uneasy sense that there is some shift between what is meant by 'belief' in these two possible ways of forming a belief, the shift can be explained by the complexity of our natural resources for believing, and the uneasy feeling can perhaps be traced to our instinctively simplified self-models, according to which there should

always be a clean and transparent answer to the question of what we believe. In fact, the psychological reality of belief may be more complex still: I have alluded to a distinction between verbatim and gist representation, but in fact we seem to encode several levels of gist on encountering a situation or narrative, alongside any verbatim representations, and there can be subtle dissociations between these various levels (for a review, see Reyna, Corbin et al. 2016). Which is the optimal way to think? Here again there is no simple answer: in practice, given the limits on our capacities to reason accurately, each of our ways of believing has a complex array of strengths and weaknesses. It seems advantageous for time-pressured and resource-limited creatures like us to have a capacity to switch between these ways of thinking, but we are a long way from understanding the principles which would best guide these switches.

7. Kyburg's sunrise, and the problem of avoiding skepticism

When Kyburg encouraged us to see high probability as sufficient for acceptance, he argued that the alternative was a skepticism in which we would be unable to accept that the sun would rise tomorrow. He placed the proposition that the sun will rise (or that one's lottery ticket will lose) in a basket alongside "*any* statistical hypothesis of the sort that everyone simply accepts" (1961, 197); in his view, such propositions are accepted rationally, despite having "nothing but their probability to recommend them" (1961, 87). According to Kyburg, because "anyone has enough evidence to render it highly probable that the sun will rise tomorrow", this is a statement that everyone should accept (1961, 82). Indeed, he sees the intuition that this belief is rational as a key source of motivation for his system: "if someone maintains that the sun will not rise tomorrow, we want to be able to say that his belief is *irrational*" (1961, 82). The reliance on intuitive motivation here makes it more awkward for Kyburg to advance his proposed solution to the lottery paradox later: in his own treatment of the lottery paradox, Kyburg concluded that we need to place restrictions on the practice of conjoining rational beliefs, noting in passing that this is "one of the most curious things about the system" (1961, 196). It seems worse than curious: Kyburg may well be right that we will want to accuse someone of being irrational if they will not agree that the sun will rise tomorrow,[5] but we arguably have an even stronger desire to accuse someone of irrationality if they accept that *p*, accept that *q*, and refuse to accept that *p* and *q*.

---

[5] Or that an apparent sunrise will happen somewhere on earth tomorrow; we set aside complications involving the heliocentric solar system, circumpolar latitudes, etc.

Kyburg is interestingly neutral on questions of cognitive procedure. In arguing that "everyone simply accepts" what he calls "statistical hypotheses", he was not suggesting that we had to be following a special style of precise statistical reasoning in doing so. He does not mandate a rule of explicit reasoning that obliges us to "round upwards" and accept any proposition one sees as highly probable. As Gregory Wheeler has emphasized in his (2007) review, what Kyburg sees as making it rational for a subject to accept that the sun will rise tomorrow is just the proposition's high objective probability given the subject's evidence; it is not that the subject is rational in virtue of first assigning this proposition a high subjective probability, and then reasoning that the proposition is acceptable on this basis. Indeed, Kyburg focuses consistently on what we would now call propositional as opposed to doxastic justification. His account is neutral on the manner in which a person must be thinking in order to count as rational: it is a substantive, rather than procedural, account of rationality. If we now think that a full account of rational acceptance will need to assess the way in which a proposition is accepted—so that it is possible to embrace a tautology irrationally, for example—then we will need to enrich Kyburg's account. At the same time, a closer look at the foundation of Kyburg's epistemology may give us fresh ideas about the avoidance of skepticism.

Kyburg introduces the expression 'rational corpus' for the set of beliefs an individual has at any given time, where it is assumed to be transparent to an agent what he believes. Membership in this corpus could in principle be settled for any individual, Kyburg says, "by presenting him with a list of statements … and asking him to check off the ones that he believes." (Kyburg 1961, 83) Whether any such belief is rational depends on its relationship to the 'basis' of the corpus, which Kyburg describes as a set of "beliefs which are based on the 'immediate experience' of the individual or group whose rational corpus we are discussing", and then glosses immediately afterwards as "those statements of matters of fact (whatever kinds of statements these may be) that are known without inference to the bearer of the rational corpus under discussion." (Kyburg 1961, 82) Any member of the basis set of noninferential knowledge is rationally believed, and any further belief is rational or irrational as a function of its probability given this basis, with high probability sufficient for rationality.

The shift from *beliefs based on immediate experience* to *matters of fact that are known without inference* is more controversial now than it was in 1961. It is not universally granted that beliefs based on immediate experience necessarily constitute knowledge, nor that they are even guaranteed to be true. The other direction of the claimed equivalence has also been challenged: it is

not obvious that what is known without inference must always be based on immediate experience. Certainly, if we mean person-level explicit inference, the type of inference that distinguishes reflective from intuitive cognition, then most epistemologists take many facts that go beyond immediate experience to be known without inference, including, for example, the identity of the current pope and the fact that the sun will rise tomorrow. Do our judgments on these matters have anything more than their probability to recommend them, or to confer a positive epistemic status upon them? Kyburg saw judgments of immediate experience as the necessary foundation of all empirical justification, but here again many contemporary epistemologists will disagree, not least out of fears that such a way of thinking could lead us to a skeptical phenomenalism of the present moment. In particular, externalists could point to features of the environmental history of the judgment as contributing to its positive epistemic status: for example, that one's judgment about the identity of the pope originally came from a knowledgeable source, and that it has been well preserved in memory in a relevantly stable world, and so on. If externalists are right that subjects can have noninferential knowledge of such matters as the identity of world leaders, the factors in virtue of which they have this knowledge are not guaranteed to be first-person accessible or retrievable by a subject who is challenged to give an argumentative defense of her judgment. The belief that the sun will rise tomorrow can have more to recommend it than the subject can state, not to mention more in its favor, epistemically, than its probability on the subject's immediate experience.

With a more expansive notion of what is known without inference (not restricted to immediate experience), we could accept that rationality was a function of the relationship to this basis without courting skepticism. At the extreme, one could hold a position in which only what is known is fully rational, allowing just the basis of what is known without inference and further additions of inferential knowledge gained by sound precise reasoning from that basis. On this way of thinking, it would make no sense to characterize "the sun will rise tomorrow" as a "statistical hypothesis." We could certainly accept propositions on the basis of sound statistical reasoning, but these would be propositions like "this ticket has a .999999 chance of winning", with the precise character of the statistical reasoning appropriately leaving its mark on the conclusion, and this conclusion believed outright. Less radical views are also possible, however. One might grant the full rationality of what is known without inference, and also allow that the stock of what it is rational to believe could then be enlarged not only precise reasoning, but also by various manipulations of the gist of what is known. To the extent that gist representations approximate verbatim originals, this proposal could maintain the idea that rational belief is best understood in terms of a relationship to

knowledge. There are limits to the rational employment of gist in governing action and reasoning: it can be irrational to rely on gist in making insurance decisions or stringing together multiple conjuncts, but it does not follow that epistemic rationality generally forbids the employment of gist representations. Rougher gist representations can stray further from knowledge than the sharp outright beliefs we gain from reasoning precisely from what is known, but in virtue of their role in guiding thought and action, they can still count as beliefs, and in virtue of their role in keeping limited creatures like us in touch with the truth, they can still count as rational.

As epistemologists, we are chiefly concerned with normative questions, but in my view a deeper descriptive understanding of the nature of belief can serve as a valuable guide in decoding the significance of our epistemic intuitions, and the appropriateness of various possible norms. Here I hope to have made some case for the relevance of our psychological complexity to the problem of the lottery paradox. For a creature with only one way of believing, the question of whether it is rational to believe that one's lottery ticket will lose might have a simple answer: for us, this is a surprisingly and revealingly interesting question.

References:

Baddeley, A. (2007). Working Memory, Thought, and Action. New York, Oxford University Press.
Britt, M. A., C. Kurby, S. Dandotkar and C. Wolfe (2007). "I agreed with what? Memory for simple argument claims." Discourse Processes **45**(1): 52-84.
Carruthers, P. (2011). The Opacity of Mind: An Integrative Theory of Self-Knowledge. New York, Oxford University Press.
Carruthers, P. (2015). The Centered Mind: What the Science of Working Memory Shows us about the Nature of Human Thought. Oxford, Oxford University Press.
Christensen, D. (2010). "Higher Order Evidence." Philosophy and Phenomenological Research **81**(1): 185-215.
Davies, M. (2008-). The Corpus of Contemporary American English: 450 million words, 1990-present.
Davies, M. and D. Gardner (2010). Frequency Dictionary of American English. New York, Routledge.
Dehé, N. and A. Wichmann (2010). "Sentence initial *I think (that)* and *I believe (that)*." Studies in Language **34**(1): 36-74.
Dodd, D. (2017). "Belief and certainty." Synthese **194**(11): 4597-4621.
Dorst, K. (2018). "Lockeans Maximize Expected Accuracy." Mind **online early**: 1-37.
Douven, I. (2008). "Knowledge and practical reasoning." Dialectica **62**(1): 101-118.
Douven, I. (2010). "The pragmatics of belief." Journal of Pragmatics **42**(1): 35-47.
Douven, I. (2012). "The lottery paradox and the pragmatics of belief." dialectica **66**(3): 351-373.
Douven, I. and T. Williamson (2006). "Generalizing the lottery paradox." The British journal for the philosophy of science **57**(4): 755-779.
Ebert, P. A., M. Smith and I. Durbach (2018). "Lottery judgments: A philosophical and experimental study." Philosophical Psychology **31**(1): 110-138.
Evans, J. (2007). Hypothetical thinking: dual processes in reasoning and judgement. New York, Psychology Press.
Evans, J. (2009). "Introspection, confabulation, and dual-process theory." Behavioral and Brain Sciences **32**(02): 142-143.

Evans, J. and K. Stanovich (2013). "Dual-Process Theories of Higher Cognition Advancing the Debate." Perspectives on Psychological Science **8**(3): 223-241.

Gigerenzer, G. and H. Brighton (2009). "Homo heuristicus: Why biased minds make better inferences." Topics in Cognitive Science **1**(1): 107-143.

Goldsmith, M., A. Koriat and A. Weinberg-Eliezer (2002). "Strategic regulation of grain size memory reporting." Journal of Experimental Psychology: General **131**(1): 73.

Hawthorne, J., D. Rothschild and L. Spectre (2016). "Belief is Weak." Philosophical Studies **173**: 1393-1404.

Kaltenböck, G. (2009). "Initial *I think*: main or comment clause?" Discourse and Interaction **2**(1): 49-70.

Kaltenböck, G. (2010). Pragmatic functions of parenthetical *I think*. New Approaches to Hedging. G. Kaltenböck, W. Mihatsch and S. Schneider. Bingley, UK, Emerald Group**:** 243-270.

Kärkkäinen, E. (2010). Position and scope of epistemic phrases in planned and unplanned American English. New Approaches to Hedging. G. Kaltenböck, Emerald Group.

Kyburg, H. (1961). Probability and the logic of rational belief. Middletown, CT, Wesleyan University Press.

Mercier, H. and D. Sperber (2017). The Enigma of Reason. Cambridge, MA, Harvard University Press.

Nagel, J. (2011). "The psychological basis of the Harman-Vogel paradox." Philosophers' Imprint **11**(5): 1-28.

Newell, B. and M. Lee (2011). "The Right Tool for the Job? Evidence Accumulation in Decision Making." Psychological review **24**(5): 456-481.

Nyarko, Y. and A. Schotter (2002). "An experimental study of belief learning using elicited beliefs." Econometrica **70**(3): 971-1005.

Ranger, G. (2018). Discourse Markers: An Enunciative Approach. Basingstoke, UK, Palgrave Macmillan.

Reyna, V. F. (2004). "How people make decisions that involve risk: A dual-processes approach." Current directions in psychological science **13**(2): 60-66.

Reyna, V. F., J. C. Corbin, R. B. Weldon and C. J. Brainerd (2016). "How fuzzy-trace theory predicts true and false memories for words, sentences, and narratives." Journal of applied research in memory and cognition **5**(1): 1-9.

Reyna, V. F., F. J. Lloyd and C. J. Brainerd (2003). Memory, development, and rationality: An integrative theory of judgment and decision making. Emerging perspectives on judgment and decision research. S. Scheider and J. Shanteau. New York, Cambridge University Press**:** 201-245.

Schlag, K. H., J. Tremewan and J. J. Van der Weele (2015). "A penny for your thoughts: A survey of methods for eliciting beliefs." Experimental Economics **18**(3): 457-490.

Simons, M. (2007). "Observations on embedding verbs, evidentiality, and presupposition." Lingua **117**(6): 1034-1056.

Sloman, S. (1996). "The empirical case for two systems of reasoning." Psychological Bulletin **119**(1): 3-22.

Slovic, P., B. Fischhoff and S. Lichtenstein (1978). "Accident probabilities and seat belt usage: A psychological perspective."

Slovic, P., B. Fischhoff, S. Lichtenstein, B. Corrigan and B. Combs (1977). "Preference for insuring against probable small losses: Insurance implications." Journal of Risk and insurance: 237-258.

Staffel, J. (2016). "Beliefs, buses and lotteries: Why rational belief can't be stably high credence." Philosophical Studies **173**(7): 1721-1734.

Stanovich, K. (2011). Rationality and the reflective mind. New York, Oxford Univ Pr.

Stanovich, K. and R. West (2008). "On the relative independence of thinking biases and cognitive ability." Journal of personality and social psychology **94**(4): 672.

Tomasello, M. (2018). "How children come to understand false beliefs: A shared intentionality account." Proceedings of the National Academy of Sciences: 201804761.

Turri, J. and O. Friedman (2014). Winners and losers in the folk epistemology of lotteries. Advances in experimental epistemology. J. R. Beebe. London, Bloomsbury**:** 45-69.

Van Bogaert, J. (2011). "*I think* and other complement-taking mental predicates: A case of and for constructional grammaticalization." Linguistics **49**(2): 295-332.

Vogel, J. (1990). Are there counterexamples to the closure principle? Doubting: Contemporary Perspectives on Skepticism. M. Roth and G. Ross. Dordrecht, Kluwer**:** 13-27.

Weisberg, J. (2013). "Knowledge in Action." Philosophers' Imprint **13**(22): 1-23.

Weisberg, J. (forthcoming). "Belief in psyontology." Philosophers' Imprint.

Wheeler, G. (2007). A review of the lottery paradox. Probability and inference: Essays in honour of Henry E. Kyburg, Jr. W. Harper and G. Wheeler. London, College Publications**:** 1-31.

WHO. (2016). "Life Tables by Country: Italy." <u>Global Health Observatory Data Repository</u> Retrieved July 11, 2018, from <u>http://apps.who.int/gho/data/?theme=main&vid=60800</u>.

Williamson, T. (2000). <u>Knowledge and its Limits</u>. New York, Oxford University Press.

Windschitl, P. D. and G. L. Wells (1996). "Measuring psychological uncertainty: Verbal versus numeric methods." <u>Journal of Experimental Psychology: Applied</u> **2**(4): 343-364.