

Basic-Know And Super-Know¹

ANNA MAHTANI

London School of Economics and Political Science

Sometimes a proposition is ‘opaque’ to an agent: (s)he doesn’t know it, but (s)he does know something about how coming to know it should affect his or her credence function. It is tempting to assume that a rational agent’s credence function coheres in a certain way with his or her knowledge of these opaque propositions, and I call this the ‘Opaque Proposition Principle’. The principle is compelling but demonstrably false. I explain this incongruity by showing that the principle is ambiguous: the term ‘know’ as it appears in the principle can be interpreted in two different ways, as either basic-know or super-know. I use this distinction to construct a plausible version of the principle, and then to similarly construct plausible versions of the Reflection Principle and the Sure-Thing Principle.

1. The Opaque Proposition Principle

I begin by illustrating the Opaque Proposition Principle with an example.

Suppose that you have bought a ticket in the World’s Smallest Lottery. There are four tickets (tickets 1, 2, 3 and 4) and four players, each of whom has bought a single ticket. There are two prizes, and so two of the four tickets have been selected randomly as the winners. Your ticket is number 1. Assuming that you are rational, what is your credence that (WIN1) your ticket number 1 is a winner?

Of course, your credence in WIN1 is $\frac{1}{2}$.² Now consider what your credence in WIN1 would be were you to learn simply that (WIN2) ticket 2 is a winning ticket. We can assume that your credence in WIN1 in this counterfactual situation would — and rationally should — equal your actual conditional credence in WIN1 given WIN2, which is $\frac{1}{3}$.³

¹ Many thanks to the Leverhulme Trust who funded this research. I am also very grateful for questions and comments from Jonathan Birch, Darren Bradley, Patrick Greenough, Bernhard Salow, the audience at a LanCog workshop on ‘Puzzles of Knowledge’ at the University of Lisbon, and a meeting of the British Society for the Philosophy of Science at the LSE.

² To check this, let’s fill in a bit of detail about the process by which the organizer picks out two tickets. First (s)he picks out one ticket at random from amongst the four; then (s)he picks out a further winning ticket from the remaining three tickets. Thus there are two ways that ticket 1 might be selected: one way is by being selected on the first draw, and the other is by not being selected on the first draw but being selected on the subsequent draw. Thus $\text{Cr}(\text{WIN1}) = \text{Cr}(\text{ticket 1 is selected on the first draw}) + \text{Cr}(\text{ticket 1 is not selected on the first draw but is selected on the subsequent draw}) = \frac{1}{4} + (\frac{3}{4} \times \frac{1}{3}) = \frac{1}{2}$.

³ $\text{Cr}(\text{WIN1}|\text{WIN2}) = \text{Cr}(\text{WIN1}\&\text{WIN2})/\text{Cr}(\text{WIN2})$. To calculate $\text{Cr}(\text{WIN1}\&\text{WIN2})$, consider that there are two ways that (WIN1&WIN2) could be true: ticket 1 could be selected on the first draw, and then ticket 2 on the second draw, or vice-versa. Thus $\text{Cr}(\text{WIN1}\&\text{WIN2}) = (\frac{1}{4} \times \frac{1}{3}) + (\frac{1}{4} \times \frac{1}{3}) = \frac{1}{6}$. $\text{Cr}(\text{WIN2}) = \frac{1}{2}$. Thus $\text{Cr}(\text{WIN1}|\text{WIN2}) = \frac{1}{3}$.

Similarly, your credence in WIN1 would be $\frac{1}{3}$ if instead you were to learn just WIN3, or if you were to learn just WIN4. As things stand, you do not know any of WIN2, WIN3 or WIN4, but you do know that at least one of these claims must be true, because there must be two winning tickets. You know then that there is *some* true proposition (either WIN2, WIN3 or WIN4) such that if you were to come to know it, you would — and rationally should — have a credence of $\frac{1}{3}$ in WIN1. Doesn't it follow that your actual current credence in WIN1 should be $\frac{1}{3}$? After all, you know that there is a true proposition that would rightly drive your credence in WIN1 down to $\frac{1}{3}$ if you but knew it. You do not know this proposition — it is *opaque* to you — but isn't knowing that there *is* such a proposition enough?⁴

Underlying this piece of reasoning is what I am calling the 'Opaque Proposition Principle', which is as follows:

If: for some claim H and value v , an agent knows that there is a true proposition E such that if (s)he were to come to know E (and not to learn or forget anything else)⁵, then his or her credence in H should be v ,

Then: that agent ought to have a credence of v in H.

This Opaque Proposition (OP) Principle is related to the principle known as 'Expert Deference' (Gaifman, 1988; Hall, 1994; Elga A., 2007; Titelbaum, 2013; Hedden, 2015), which in turn is an adaptation of the Reflection Principle (van Fraassen, 1984). Expert Deference states that for any agent that you recognize as an expert, you (with credence function Cr_{you}) ought to defer to that expert (with credence function Cr_{ex}), in the following sense: for any claim H and value v such that $Cr_{you}(Cr_{ex}(H)=v) > 0$, $Cr(H|Cr_{ex}(H)=v) = v$. The definition of an expert on this view is an agent who shares your priors, knows all that you know plus some additional evidence (call it E), and is perfectly rational and so has simply conditionalized on E.⁶ The OP Principle is in one sense more general than Expert Deference, because the OP Principle requires you to defer to the credence function that you *would* (and should) have *were* you to conditionalize on E — i.e. to defer to a merely counterfactual credence function — whereas Expert Deference applies only when there is some actual individual, either yourself at a later time, or some other person, who actually has conditionalized on E. In another sense the OP Principle is narrower than Expert Deference, as the OP Principle applies only when you *know* what credence you should have in H were you to conditionalize on E. We can generalize the OP Principle, but in doing so we brush over the ambiguity that I wish to highlight. Here then I will briefly mention the generalized formalization, before returning to our original version.

⁴ The claims WIN2, WIN3 and WIN4 are not disjoint, for it may be, for example, that WIN2 and WIN3 are both true. The claim that there is *some* true proposition E amongst these can be read in two different ways: it can be read as the claim that there is at least one true proposition amongst these; or it can be read as the claim that there is exactly one true proposition (namely E) amongst these. The focus of section 3 below is a more careful analysis of this sort of claim and its ambiguities.

⁵ It might here be objected that whenever an agent learns some proposition 'E', he or she will inevitably learn in addition *how* (s)he came to learn E. This is not a necessary truth: it is at the very least *logically* possible for an agent to learn some proposition E without also learning anything about how (s)he came to learn E. This is all that is needed for the OP principle to be coherent. I delve more deeply into this issue in 4 (on 'Selection Bias').

⁶ Of course, the notion of an 'expert' is relative to some agent at a time. Someone may be an expert relative to me now (i.e. share my priors, know all I currently know and more), but not be an expert relative to you.

For the generalized formalization, let Cr designate the credence function of an agent, and Cr_E designate the credence function that that agent would have were (s)he to come to know just E and conditionalize, as rationality demands, on what (s)he has learnt. Then we require that for any claim H and value v such that $Cr(Cr_E(H)=v) > 0$, $Cr(H|Cr_E(H)=v) = v$. We will return to this generalized formulation later: for now, we focus on the original version of the OP Principle.⁷

The OP Principle is compelling at first sight. It is definitely wrong however. To demonstrate this, I now give an argument parallel to that given above, but yielding a different and incompatible conclusion. First consider how your credence in WIN1 would (and should) change were you to learn that (LOSE2) ticket 2 is a losing ticket. This is equivalent to your current conditional credence in WIN1 given LOSE2, which is $\frac{2}{3}$.⁸ Similarly, your credence in WIN1 would be $\frac{2}{3}$ were you to learn LOSE3, or were you to learn LOSE4. As things stand, you know that at least one of LOSE2, LOSE3 or LOSE4 must be true (because only 2 tickets win), and so you know that there is *some* proposition such that if you were to come to know it, you should have a credence of $\frac{2}{3}$ in WIN1. Thus by the OP Principle it follows that your credence in WIN1 should be $\frac{2}{3}$. This contradicts our earlier conclusion that your credence in WIN1 should be $\frac{1}{3}$: your credence in WIN1 cannot be required to be both $\frac{2}{3}$ and $\frac{1}{3}$.

Thus the OP Principle is certainly wrong, and yet it has an intuitive appeal. My aim in this paper is to reconcile these two facts. I argue that the principle is correct under one reading, but incorrect under another. To distinguish these two readings, I first need to tease apart two different senses of the word ‘know’.

2. Two Senses of Know

There are many different grammatical constructions involving the word ‘know’ and its derivations. For example, we might say, ‘Kate knows Sam’: here ‘knows’ is followed by a name, and we are attributing knowledge-by-acquaintance to Kate. Or we might say, ‘Kate knows how to tie her shoelaces’, and here ‘knows’ is followed by a ‘how to’ clause, and we are attributing something like an ability to Kate. Or we might say, ‘Kate knows that the Battle of Hastings happened in 1066’, and here ‘knows’ is followed by a ‘that’ clause, and we are attributing propositional knowledge to Kate. If we now let E denote the proposition that the Battle of Hastings happened in 1066, then we can say, ‘Kate knows E ’. Here ‘knows’ is followed by a name, but we are not attributing knowledge-by-acquaintance to Kate: E denotes a proposition, and we are attributing propositional knowledge to Kate. This is the sort of construction that I am interested in: where

⁷ We could also consider how the OP principle relates to Feldman’s principle that ‘evidence for evidence is evidence’ (the ‘EEE’ principle) (Feldman, 2007). Though the principles are not the same, and neither entails the other, there are nevertheless interesting relations between them. Of particular interest are attempts to disambiguate the EEE principle (Fitelson, 2012), and the discussion in the literature over whether the EEE principle should be given a *de re* or *de dicto* reading (Roche, 2014; Tal & Comesaña, 2015) which bears interesting relations to my discussion in section 3 over whether E should be taken as a variable or a name.

⁸ $Cr(WIN1|LOSE2) = Cr(WIN1 \& LOSE2) / Cr(LOSE2)$. To calculate $Cr(WIN1 \& LOSE2)$, consider that there are four ways that (WIN1&LOSE2) could be true: ticket 1 is drawn first and ticket 3 drawn second; ticket 1 is drawn first and ticket 4 drawn second; ticket 3 is drawn first and ticket 1 drawn second; ticket 4 is drawn first and ticket 1 drawn second. Thus $Cr(WIN1 \& LOSE2) = (\frac{1}{4} \times \frac{1}{3}) + (\frac{1}{4} \times \frac{1}{3}) + (\frac{1}{4} \times \frac{1}{3}) + (\frac{1}{4} \times \frac{1}{3}) = \frac{4}{12} = \frac{1}{3}$. $Cr(LOSE2) = \frac{1}{2}$. Thus $Cr(WIN1|LOSE2) = \frac{2}{3}$.

‘knows’ is followed by a name or definite description, and where that name or definite description denotes a proposition, and what is attributed is propositional knowledge.⁹

I claim that there are two ways of reading ‘know’ when it appears in this sort of construction, and I call these ‘basic-know’ and ‘super-know’. To basic-know E is simply to know the proposition that ‘E’ denotes. To super-know E is to know *which* proposition ‘E’ denotes. It is possible then to basic-know E without super-knowing E: you might know E, without knowing that that proposition is E. It is also possible to super-know E without basic-knowing E: you might know which proposition E is, without knowing that proposition — perhaps because you don’t believe it.

Here is an example to illustrate these two different readings of ‘know’.

The Guru’s Second Tenet

You are in the process of joining a religious organization. The leaders want to figure out whether you are worthy. They wonder — do you know any of the Guru’s Twelve Tenets? Every member of the organization is expected to know at least one of these. This is the first time that you have heard of the Guru and his Tenets, and you are quite unable to tell the leaders what any of them are. The leaders conclude that as you do not know any of the tenets, you need to attend an introductory course before you can be considered a member.

Now as it happens, the Guru’s Second Tenet is that knowledge is the greatest gift, and through your own insight or experience you do already know that knowledge is the greatest gift — you just don’t know that this is the Guru’s Second Tenet. In a sense then you already know the Guru’s Second Tenet: you basic-know it, but you don’t super-know it. Have the leaders made a mistake in concluding that you don’t know any of the Guru’s Twelve Tenets? It’s not obvious: perhaps it depends on which sense of ‘know’ (basic-know or super-know) the leaders have in mind, or would have in mind if they considered the distinction.

We can contrast your state of mind with that of another candidate who can confidently identify each of the Guru’s Twelve Tenets with a proposition, but does not know any of them. For example, when asked what the Guru’s Second Tenet is, this candidate correctly identifies this as the proposition that knowledge is the greatest gift, but (s)he does not actually believe — let alone know — that knowledge is the greatest gift. Does this candidate know the Guru’s Second Tenet? (S)he knows it in one sense, in that (s)he super-knows it, but (s)he does not basic-know it. Whether (s)he should qualify as knowing this tenet depends on what sense of ‘know’ is appropriate here.

This scenario of the Guru’s Twelve Tenets illustrates the possible readings of ‘know’ as it appears in the construction that I am interested in. There are numerous other scenarios we could explore along the same lines. For example, for a pupil at a school to count as knowing Pythagoras’ Theorem, does (s)he need to basic-know it, super-know it, or both? When the doctor asks whether you know the conclusion of the latest study on back-pain, is (s)he asking whether you basic-know it, or super-know it? Wherever the construction that I am interested in occurs, we can ask how ‘know’ should be read. Which reading is correct

⁹ Some people think that propositions are too coarse-grained to be the objects of knowledge: see (King, Soames, & Speaks, 2014) for the history and current thinking on this issue. My point here can be reworded to accommodate various views in this debate — e.g. we might take E to be a sentence-like entity rather than a proposition, or we might take ‘knows’ in this construction to denote a three-place relation between a person, a proposition (such as E) and a guise.

depends on the details of the scenario: in one context it may be more natural to read ‘know’ as basic-know; in another context, super-know may be the more natural reading; and in some scenarios either reading may be appropriate and there will be genuine ambiguity.

Here I add a final comment on the distinction, which will become relevant later in this paper. The term ‘super-know’ introduces a certain sort of intensionality. It is often said that psychological terms such as ‘know’ introduce an intensional context: for example many would claim that a rational person can know that Hesperus is visible in the evening, without knowing that Phosphorus is visible in the evening. Various moves can be made in response to this thought: we might claim that propositions are fine-grained enough to distinguish between these scenarios; or that the objects of belief are not (merely) propositions; and other solutions are possible. With super-knowledge, however, we introduce a further layer of intensionality. Whatever the objects of belief turn out to be, it is possible for a rational agent to super-know such an object under one designator, but not under another. This is because whether an agent super-knows some object of belief E depends on whether she recognizes that object as E . Thus for example suppose that Pythagoras Theorem and the proposition (or whatever is the correct type of thing to be an object of belief) expressed by the first sentence on page 14 of my maths book, are identical. Then if I basic-know one, I basic-know the other, but I might super-know one without super-knowing the other. For example, I might know what proposition Pythagoras’ Theorem is, without knowing what the proposition expressed by the first sentence on page 14 of my maths book is. In this way, super-knowledge (but not basic-knowledge) introduces a further layer of intensionality.

Having teased apart these two senses of the word ‘know’, I turn to apply this distinction to the puzzle with which we began.

3. Disambiguating the Opaque Proposition Principle

We began with the OP Principle, which is intuitively compelling but demonstrably false:

If: for some claim H and value v , an agent knows that there is a true proposition E such that if (s)he were to come to know* E (and not to learn or forget anything else) then his or her credence in H should be v ,

Then: that agent ought to have a credence of v in H .

The word ‘know’ (or derivations of it) appears twice in this piece of reasoning. In its first appearance, it is followed by a ‘that’ clause, so this is not the construction that interests us. In its second appearance (where it is starred) it is followed by ‘ E ’, which looks like a name for a proposition. This is an example of the construction that I am interested in, and here there are two different ways of taking the word ‘know’: as basic-know or as super-know.¹⁰ I claim that the OP Principle is true when we take ‘know’ here to mean super-know, but false when we take it to mean basic-know. This explains why the principle is both intuitively compelling but demonstrably false: the principle is intuitively compelling because it is actually true under one reading, even though under a different reading the principle is false.

I begin by showing that the OP Principle is false when the relevant instance of ‘know’ is taken to mean basic-know.

¹⁰ Note that the antecedent entails that the agent knows that E is true, so if the agent super-knows E , she also knows that E is true, and thereby basic-knows E .

The 'basic-know' reading

Here we face yet more complexity, for if we take the relevant instance of 'know' to mean basic-know, then there are at least two ways of interpreting the antecedent clause of the OP Principle. I wrote above that 'E' looks like a name for a proposition, but 'E' could alternatively be taken as a variable ranging over propositions.¹¹ I show that on either of these interpretations (i.e. regardless of whether we take 'E' as a name or as a variable), provided that we take 'know' to mean basic-know, the OP Principle is false.

We can see this by returning to our case of the World's Smallest Lottery. Recall that you hold ticket 1, and clearly your credence that (WIN1) your ticket will win should be $\frac{1}{2}$. Yet you know that at least one of the claims WIN2, WIN3 and WIN4 must be true, and that if you were to come to know any one of these claims then your credence in WIN1 should become $\frac{1}{3}$. Interpreting 'know' as basic-know, and taking 'E' as a variable, then, we can see then that the following claim is true in our world's smallest lottery case:

You know that there exists some true proposition E such that if you were to come to basic-know E (and not to learn or forget anything else) then your credence in WIN1 should be $\frac{1}{3}$.¹²

This claim (including the whole conditional embedded within it) is an instantiation of the antecedent of the OP principle. So — if we allow the OP principle to guide us here — the relevant instantiation of the consequent of the OP principle follows: your credence in WIN1 should be $\frac{1}{3}$. But of course it isn't the case that your credence in WIN1 should be $\frac{1}{3}$. Thus we can see that the OP Principle is false when we take the relevant instance of 'know' to mean 'basic-know', and 'E' as a variable.¹³

What happens if we continue to read 'know' as basic-know, but rather than take E as a variable, we take it as a name denoting some particular proposition? To illustrate this reading, let us return to the World's Smallest Lottery example, and give a definite definition to fix the referent of the name 'E'.¹⁴ First let us define the predicate G as follows: a proposition is G iff it meets the following criteria:

- i) If WIN2 is true, then WIN2 is G (and no other proposition is)
- ii) If WIN2 is false and WIN3 is true, then WIN3 is G (and no other proposition is)
- iii) If both WIN2 and WIN3 are false, then WIN4 is G (and no other proposition is)

¹¹ The principle is written in English (albeit rather technical English), so there is no rule of syntax to determine whether 'E' is to be read as a name or variable.

¹² Of course, there may be more than one such proposition. For example, it may be that WIN2 and WIN3 are both true, and coming to basic-know either of these propositions would rightly drive your credence in WIN1 down to $\frac{1}{3}$. Thus the claim is not that you know that there exists *exactly one* true proposition E such that if you were to come to basic-know E then your credence in WIN1 should be $\frac{1}{3}$, but rather than you know that there exists *at least one* such true proposition.

¹³ As before, both here and below we could show that the OP Principle so interpreted leads not just to a counterintuitive conclusion, but to contradictory conclusions.

¹⁴ There are many other ways that 'E' could be defined: here I just pick one way, as this is sufficient to show that the OP principle (with the relevant instance of 'know' read as basic-know) is false.

Now we can use a definite description — ‘the proposition that is G’ — to denote a proposition, and we can stipulate that ‘E’ denotes the proposition so described.¹⁵ Note that this description will denote one and only one proposition, and the proposition so denoted will be true — for it will denote WIN2 only if WIN2 is true, WIN3 only if WIN3 is true (and WIN2 is false), and WIN4 only if both WIN2 and WIN3 are false in which case WIN4 will be true (because at least two tickets must win).¹⁶ We can now consider the following claim, which is an instantiation of the antecedent of the OP principle, with the second instance of ‘know’ read as basic-know:

You know that there is a true proposition, namely E (defined as above), such that if you were to come to basic-know E (and not to learn or forget anything else), then your credence in WIN1 should be $\frac{1}{3}$.

This claim is true. For you know that E will be either WIN2, WIN3, or WIN4. And so you know that if you were to come to basic-know E, you would just come to (basic-) know either WIN2, WIN3, or WIN4.¹⁷ You know that if you were to come to (basic-) know WIN2 (without learning or forgetting anything else) then your credence in WIN1 should be $\frac{1}{3}$ — and the same holds if you were to come to (basic-)know either WIN3 or WIN4.¹⁸ Thus this instantiation of the antecedent of the OP principle is true, and so if the OP principle holds then the relevant instantiation of the consequent follows: you ought to have a credence of $\frac{1}{3}$ in WIN1. But it is *not* the case that you ought to have a credence of $\frac{1}{3}$ in WIN1, and so the OP principle leads us astray here. Thus the OP principle is false if we take the relevant instance of ‘know’ to mean basic-know, and read ‘E’ as a proper name.

Thus I have shown that — whether we take ‘E’ as a name or a variable — the OP Principle is false when the relevant instance of ‘know’ is read as basic-know. I turn now to consider how the OP Principle fares when the relevant instance of ‘know’ is read as super-know.

The ‘super-know’ reading

On this reading, ‘E’ must be interpreted as a name for a proposition, rather than as a variable. This is because to super-know E is to know which proposition ‘E’ denotes, and if

¹⁵ I do not here imply that every name is merely a disguised definite description. But my argument does depend on there being *some* names which are ‘descriptive names’ — i.e. ‘names whose referent is fixed by a description’ (Evans, 1979, p. 162). For in order to show that the OP principle is false when the relevant instance of ‘know’ is taken to mean basic-know, and ‘E’ is taken as a name, the instantiation of the principle that I use as a counterexample in the text interprets ‘E’ as a descriptive name.

I should add that we can rigidify the definite description which fixes the referent of E without affecting my argument in this paper.

¹⁶ The name E denotes some proposition, or event. But you do not know which proposition (out of WIN2, WIN3 and WIN4) E denotes: it is *opaque*. We can and do reason about propositions denoted in this sort of way. For example, we might learn that *what Bob said* wasn’t very kind, and conditionalize on this piece of information, without knowing what Bob said.

¹⁷ WIN2, WIN3 and WIN4 are not disjoint, for it may be that two of these propositions are true together. Thus the set of propositions that you think you might come to know when you come to basic-know E do not form a partition.

¹⁸ Here again I am assuming (reasonably, in my view) that it is possible for an agent to come to (basic-) know some proposition (e.g. WIN2) without also coming to know how (s)he came to know it. I engage with an objection to this view in section 4 (on ‘Selection Bias’).

‘E’ is a variable then ‘E’ does not denote at all, and so what it would be to super-know E is undefined. Let us suppose then as before that in the case of the World’s Smallest Lottery, ‘E’ names the proposition which is G. Let us now consider the following claim:

You know that there is a proposition, namely E (defined as above), such that if you were to come to super-know E, then your credence in WIN1 should be $\frac{1}{3}$.

This claim (including its embedded conditional) is an instantiation of the antecedent of the OP Principle (with the relevant instance of ‘know’ read as super-know). If this claim were true, then the OP principle would lead us to the false conclusion that your credence in WIN1 ought to be $\frac{1}{3}$. But fortunately, the claim is not true. To see this, consider that you know that E is WIN2, WIN3 or WIN4, but you don’t know which it is. Let us go through each possibility in turn. Firstly, you know that if you were to come to super-know E, and E was WIN2, then your credence in WIN1 should indeed be $\frac{1}{3}$. This is because to come to super-know E would be to come to know that E is WIN2, and so that WIN2 is G, and from this you can (only) deduce that WIN2 is true. Secondly, you know that if you were to come to super-know E, and E was WIN3, then your credence in WIN1 should remain at $\frac{1}{2}$. This is because to come to super-know E would be to come to know that E is WIN3, and so that WIN3 is G, and from this you can deduce both that WIN2 is false and that WIN3 is true. Finally you know that if you were to come to super-know E, and E was WIN4, then your credence in WIN1 should increase to 1. This is because to come to super-know E would be to come to know that E is WIN4, and so that WIN4 is G, and from this you can deduce that both WIN2 and WIN3 must be false, and so that WIN1 and WIN4 must be true. Thus we can see that it is *not* the case that you know that if you were to come to super-know E, then your credence in WIN1 should be $\frac{1}{3}$. Indeed given the definition of E that we are working with here, there is no n such that you know that if you were to come to super-know E, then your credence in WIN1 should be n : you do not know what your credence in WIN1 ought to be — whether it ought to be $\frac{1}{3}$, $\frac{1}{2}$ or 1 — on coming to super-know E, for that will depend on what proposition E is. Thus the OP principle does not lead us into trouble here.

More generally, no matter how we define E, provided we interpret the relevant instance of ‘know’ as super-know, the OP Principle will not lead us astray. To see this, let us start by assuming the antecedent of the OP principle (with ‘know’ read as super-know), and show that the consequent can be derived. The antecedent is as follows: for some H and v , an agent knows that there is some proposition E such that if (s)he were to come to super-know E, then his or her credence in H should be v . To come to super-know E is just to come to (basic-)know some proposition of the form *E is the proposition that such-and-such*. Let $S = \{S_1, S_2, \dots, S_n\}$ be the set of propositions that the agent thinks might be what (s)he needs to *basic*-know in order to come to *super*-know E.¹⁹ More precisely, for every proposition S_i such that the agent thinks that coming to super-know E might consist in coming to basic-know S_i , S_i will be a member of the set S; and for every proposition S_i in S, the agent thinks that coming to super-know E might consist in coming to basic-know S_i ; and finally the same proposition will not appear twice in the set S. It follows that the

¹⁹ So for example in our case of the world’s smallest lottery, with E defined using a definite description as on pages 6–7, the set S of propositions that you think might be what you need to basic-know in order to come to super-know E, will be {E is WIN2, E is WIN3, E is WIN4}.

set S is a partition: at least one proposition in S is true (for E must be *some* proposition), and no more than one proposition in S is true (for E can only be *one* proposition).²⁰ Then the antecedent of the OP Principle — i.e. the claim that the agent knows that if (s)he were to come to super-know E , then his or her credence in H should be v — is equivalent to the claim that the agent knows that if (s)he were to come to basic-know any proposition in the set S , then his or her credence in H should be v . Given that the set S is a partition, it follows automatically from the probability axioms that the agent's credence in H should be v .²¹ Thus the consequent of the appealing reasoning is guaranteed to follow from the antecedent. Provided then that we interpret the relevant instance of 'know' as super-know, the appealing reasoning cannot lead us astray.^{22,23}

Thus I have shown that the OP Principle is false when the relevant instance of 'know' is interpreted as basic-know, but true when it is interpreted as super-know. I turn now to the generalized version of the OP Principle that I mentioned in section 1. Here we let Cr designate the credence function of an agent, and Cr_E designate the credence function that that agent would have were (s)he to come to know E and simply conditionalize on that new evidence, as rationality demands. Then the generalized version of the OP Principle requires that for any claim H and value v such that $Cr(Cr_E(H)=v) > 0$, $Cr(H|Cr_E(H)=v) = v$. I wrote in section 1 that this generalized version skated over the ambiguity that I wanted to highlight, and we can now pinpoint where this happens: in the definition of Cr_E . We can see that this notation involves a name 'E' for a proposition, and thus we can draw a distinction here like that between basic-know and super-know. Where Cr_E is defined as the credence function that the agent would have were (s)he to come to know E and conditionalize as rationality demands, the instance of 'know' immediately before 'E' could be taken as either basic-know or super-know. I propose for clarity that we adopt some new terminology: Cr_E is the credence function that the agent would have were she to come to basic-know E (and conditionalize as rationality demands), and Cr_{E^*} is the credence function that the agent would have were she to come to super-know E (and conditionalize as rationality demands). Then we can consider two different versions of the generalized OP Principle. The first requires that for any claim H and value v such that Cr

²⁰ That E must denote one and only one proposition follows from the fact that the agent *knows that there is some proposition E* — given that we are taking 'E' as a name rather than a variable.

²¹ To see this, consider first that knowledge is factive, so from the claim that the agent *knows* that if (s)he were to come to basic-know any proposition in the set S , then his or her credence in H should be v , it follows that if the agent were to come to basic-know any proposition in the set S , then his or her credence in H should be v . Thus it follows that for any $S_i \in S$, $Cr(H|S_i) = v$. From $Cr(H|S_i) = v$, we can infer that $Cr(H \& S_i) / Cr(S_i) = v$, and so $Cr(H \& S_i) = v Cr(S_i)$. Thus from the claim that for any $S_i \in S$, $Cr(H|S_i) = v$, we can infer that $Cr(H \& S_1) + Cr(H \& S_2) + \dots + Cr(H \& S_n) = v Cr(S_1) + v Cr(S_2) + \dots + v Cr(S_n)$. Given that S is a partition, $Cr(H \& S_1) + Cr(H \& S_2) + \dots + Cr(H \& S_n) = Cr(H)$, and $Cr(S_1) + Cr(S_2) + \dots + Cr(S_n) = 1$. Thus it follows that $Cr(H) = v$.

²² Proofs of 'Modified Reflection' and 'Expert Deference' can be found in (Weisberg, 2005), (Briggs, 2009) and (Hedden, 2015). These proofs relate in interesting ways to my proof of the OP principle here, though they rely on different assumptions.

²³ Here there is an interesting link to the Monty Hall Problem (Selvin, 1975; Savant, 1990) and I thank an anonymous reviewer for stressing this connection. When (in the Monty Hall scenario) the gameshow host opens a door revealing a goat, you do not merely come to *basic-know* the information that he has chosen to reveal to you, which might be for example that door number 2 contains a goat. Rather, you come to *super-know* the information that he has chosen to reveal to you: in addition to knowing that (say) door number 2 contains a goat, you also know that this is the information that the gameshow host has chosen to reveal to you. It is this super-knowledge, along with some background information about the set-up, which allows you to arrive at the right decision — to switch doors.

$(Cr_E(H)=v)>0$, $Cr(H|Cr_E(H)=v)=v$. The second requires that for any claim H and value v such that $Cr(Cr_{E^*}(H)=v)>0$, $Cr(H|Cr_{E^*}(H)=v)=v$. I claim — in parallel with my claims about the non-generalized OP Principle — that the first principle involving Cr_E is false, but the second principle involving Cr_{E^*} is true.

Having clarified how the OP Principles should be interpreted, I turn now to briefly consider the relevance of my points for a range of other debates: selection bias, the Reflection Principle, and the Sure-Thing Principle.

4. Selection Bias²⁴

When an agent comes to know a proposition, there is often some ‘selection process’, or ‘information-gathering process’ that led to the agent’s coming to know that proposition.²⁵ A rational agent will take into account not just the most obvious proposition learnt, but also any evidence that (s)he has about *how* (s)he came to learn that proposition.²⁶ We might think that this point in itself is enough to rescue the OP Principle, and that my distinction between basic-knowledge and super-knowledge is not needed. My response to this, in short, is to point out that sometimes an agent needs to super-know (rather than merely basic-know) a proposition, in order to be able to bring his or her evidence about the information-gathering process to bear.

I now work through this reasoning more slowly, using the case of the World’s Smallest Lottery as a guide. Recall that I used this case to argue that the OP Principle — with the relevant instance of ‘know’ interpreted as basic-know — is false. The thought was that you know that there is a proposition E (either WIN2, WIN3 or WIN4) such that if you were to come to (basic-)know E , then your credence in WIN1 would rationally go to $1/3$. By the OP Principle, it follows that your credence in WIN1 ought to be $1/3$ at once — but this is clearly false, and so, I argued, the OP Principle (with the relevant instance of ‘know’ interpreted as basic-know) leads us astray. But now we might wonder whether it is true that if you were to come to (basic-)know E , then your credence in WIN1 would rationally go to $1/3$. A rational agent ought to take into account not just the most obvious proposition learnt, but also any evidence (s)he has about the information-gathering process. If you were to take all relevant evidence into account in the counterfactual scenario where you come to (basic-)know E , then is it true that your credence in WIN1 would go to $1/3$?

In the original scenario, I specified the relevant counterfactual case as one where you come to (basic-) know *just* E . But it might be objected that there is no such counterfactual case, for surely if you were to learn, say, WIN3, you would also learn *something* about how you came to gain this information.²⁷ To deal with this objection, I flesh out

²⁴ Thanks to Darren Bradley for suggesting the relevance of this issue to my arguments.

²⁵ I say ‘often’ rather than ‘always’, because while of course there will always be processes at work that lead to an agent’s coming to know a piece of information, in many cases it is not obvious that one process in particular is the salient one.

²⁶ There is disagreement over exactly *how* this evidence of the selection/information-gathering process should be taken into account. Some recent discussion can be found in (Hutchison, 1999; Korb, 1994; Kotzen, 2012).

²⁷ I maintain that it is possible to learn some proposition without also learning something about how you came to learn this proposition (see footnote 5), and so I think that the scenario is coherent as originally described. I offer the argument in the main text here as a response to an objector who disagrees — i.e. who maintains that necessarily whenever you learn a proposition, you also learn something about how you came to learn it, and so that the scenario as originally described is incoherent.

the scenario as follows. Suppose that, having drawn the winning tickets, the organizer considers letting you know the outcome of one of tickets 2–4. (S)he decides not to do this, but easily could have, so there are close possible worlds where (s)he does. In these close possible worlds, (s)he chooses which ticket to reveal to you by random selection (so each of tickets 2–4 is equally likely to be selected), and you know that this is the set-up. Thus there is a counterfactual case in which you come to know the outcome of ticket 2, and here you know what the information-gathering process was: the organizer selected a ticket from amongst tickets 2, 3 and 4 at random, and told you the outcome for that ticket regardless of whether it won or lost. Similarly there are counterfactual cases in which you come to know the outcomes of tickets 3 and 4 in a similar way. Thus there must be a counterfactual case in which you come to know proposition E, where E is defined in the usual way, as the proposition that is G. In the counterfactual case where you come to know E, your credence in WIN1 should be $1/3$.²⁸ You know this to be the case, so this instance of the antecedent of the OP Principle holds: you know there is some proposition E such that if you were to come to (basic-)know E, your credence in WIN1 should be $1/3$. Thus the OP Principle leads us astray here, even when we are careful to remember that a rational agent takes all his or her relevant evidence into account.

The problem here is that the agent cannot bring all his or her evidence to bear in the counterfactual case, because in that case (s)he does not super-know E, and so does not know that the proposition that she has learnt *is* E. We can see E as a proposition that is ‘selected’ in a particular way from amongst the population of true propositions that express the outcome for tickets 2–4.²⁹ This selection process does not sample at random from amongst this population. Rather, it is biased towards propositions about tickets that have won; and it is biased towards propositions that concern tickets with smaller numbers. Thus E is most likely to be the proposition WIN2 (with probability $1/2$); it is next most likely to be proposition WIN3 (with probability $1/3$); then proposition WIN4 (with probability $1/6$); and it cannot be any of propositions LOSE2, LOSE3 or LOSE4. In a counterfactual case in which the agent comes to basic-know but not super-know E, the agent cannot make use of this information about how the proposition that (s)he has learnt has been selected. However, in a counterfactual case in which the agent comes to also super-know E, the agent can identify the proposition that (s)he has learnt *as* E, and so (s) he can take into account the information that (s)he has about the selection procedure. This extra information may affect the agent’s credence in WIN1, and this is how — with the relevant instance of ‘know’ taken as super-know — we manage to avoid the troublesome instantiation of the OP Principle.

In summary then I have argued that we cannot rescue the OP Principle merely by reflecting that a rational agent will take into account any evidence that (s)he has about

²⁸ To see this, suppose that E is WIN3: similar reasoning applies if E is instead either WIN2 or WIN4. Then when WIN3 is revealed to you, you come to know both WIN3 (E), and (F) the fact that ticket 3 has been randomly selected from amongst tickets 2, 3 and 4 to be revealed to you. Thus your new credence in WIN1 should equal your prior credence conditional on this new evidence, i.e. $\text{Cr}(\text{WIN1} | \text{WIN3} \& \text{F}) = \text{Cr}(\text{WIN1} \& \text{WIN3} \& \text{F}) / \text{Cr}(\text{WIN3} \& \text{F})$. $\text{Cr}(\text{WIN1} \& \text{WIN3} \& \text{F}) = \text{Cr}(\text{WIN1} \& \text{WIN3}) \times \text{Cr}(\text{F})$, because F is independent of (WIN1 & WIN3): which ticket was randomly selected to be revealed to you does not depend on which tickets have been selected to win the lottery. Thus $\text{Cr}(\text{WIN1} \& \text{WIN3} \& \text{F}) = \text{Cr}(\text{WIN1} \& \text{WIN3}) \times \text{Cr}(\text{F}) = 1/6 \times 1/3 = 1/18$. Similarly $\text{Cr}(\text{WIN3} \& \text{F}) = \text{Cr}(\text{WIN3}) \times \text{Cr}(\text{F}) = 1/2 \times 1/3 = 1/6$. Thus $\text{Cr}(\text{WIN1} | \text{WIN3} \& \text{F}) = 1/18 / 1/6 = 1/3$.

²⁹ Here ‘selected’ is metaphorical: nobody has selected the proposition. For another metaphorical use of this term, see (Bradley, 2012, pp. 155-6).

selection procedures. An agent has to recognize the proposition that (s)he has come to know as E before (s)he can make full use of any evidence (s)he has about how proposition E was selected. And an agent will recognize the proposition that (s)he has come to know as E only if (s)he super-knows that proposition. Thus the distinction between basic-knowledge and super-knowledge is necessary to rescue the OP Principle.

I turn now to apply the points that I have made to two topics in the literature: the Reflection Principle and the Sure-Thing Principle.

5. The Reflection Principle

The Reflection Principle in its original form strikes many as implausible (Christensen, 1991; Talbott, 1991), so here I focus on a much more compelling version of the principle, ‘Expert Deference’ (Gaifman, 1988; Hall, 1994; Elga A., 2007; Titelbaum, 2013; Hedden, 2015) which we can state as follows: you ought to defer to any agent whom you regard as an expert, where you (with credence function Cr) regard an agent as an expert iff you know that for some proposition E, the expert’s credence function Cr_{ex} is identical to your own credence function conditionalized on E (i.e. Cr_E).

As I mentioned in section 1, Expert Deference is closely related to the OP Principle: indeed it follows from the generalized version of the OP Principle. As such, though Expert Deference appears so plausible, it faces a number of puzzling counterexamples (Elga A., 2000; Arntzenius, 2003; Bovens & Rabinowicz, 2011; Bronfman, 2015; Mahtani, 2016). I begin by constructing a typical counterexample using my scenario of the World’s Smallest Lottery.

The scenario is as before, except this time you have brought 3 friends with you: Alice, Bob and Carol. The lottery organizer tells Alice the outcome of ticket 2, Bob the outcome of ticket 3, and Carol the outcome of ticket 4. Thus one of these three friends has learnt proposition E (where once again E is the proposition that is G). You are certain that all three friends have all the relevant evidence that you have, and have simply conditionalized on the extra true proposition that they have each learnt. Let us call the friend who has learnt proposition E the ‘E-person’. The E-person may not know that (s)he is the E-person. Nevertheless by ‘Expert Deference’ you ought to defer to this E-person, and furthermore you know that the E-person has a credence of $1/3$ in WIN1.³⁰ Thus, by Expert Deference, you too ought to have a credence of $1/3$ in WIN1. Clearly however this is not correct, and so Expert Deference has led us astray here.³¹

³⁰ To see this, suppose that E happens to be WIN3: the reasoning goes through similarly if we suppose instead that E is WIN2 or WIN4. Then Bob is the E-person, because he is the person who gets to know the outcome of ticket 3. We can assume for simplicity that the scenario is set up so that Bob knows in advance that he will get to learn the outcome of ticket 3 (just as Alice knows that she’ll get to learn the outcome of ticket 2, and Carol knows that she’ll get to learn the outcome of ticket 4). Then Bob can be sure that if WIN3 is true, then he will learn it from the organizer, and similarly if LOSE3 is true then he will learn this from the organizer. Thus on learning WIN3 he can conditionalize straightforwardly on this claim, and so will have a credence of $1/3$ in WIN1.

³¹ It is clearly not correct, because it is counterintuitive: obviously you ought to have a credence of $1/2$, not $1/3$, in WIN1. If you prefer not to rely on intuition here, then we can also show that Expert Deference entails inconsistent claims. For consider that just as one of your friends must have learnt proposition E, so another must have learnt proposition E’, defined as follows: E’ is LOSE2 if LOSE2 is true; otherwise E’ is LOSE3 if LOSE3 is true; otherwise E’ is LOSE4. By expert deference you ought to defer to this friend (the ‘E’-person’), and you can be sure that this friend has a credence of $2/3$ in WIN1, so Expert Deference requires you to have a credence of $2/3$ in WIN1. Thus Expert Deference requires you to have both a credence of $1/3$ and a credence of $2/3$ in WIN1, and this is contradictory.

The correct strategy for rescuing Expert Deference parallels the correct strategy for rescuing the generalized version of the OP Principle: we need to draw a distinction between basic-knowledge and super-knowledge. You are not required to defer to an agent who you know has come to *basic*-know E, but you are required to defer to an agent who you know has come to *super*-know E. To put this more precisely we need to make use of our new terminology: if your credence function is Cr, then your credence function updated as rationality requires on coming to basic-know E, is Cr_E , and your credence function updated as rationality requires on coming to super-know E, is Cr_{E^*} . Then we can refine our definition of an expert. You (with credence function Cr) regard an agent as an expert iff you know that for some proposition E, the expert's credence function Cr_{ex} is identical to your own credence function updated as rationality requires on coming to super-know E (i.e. Cr_{E^*}).

We can now deal with our counterexample to Group Reflection straightforwardly. The E-person — i.e. the friend who has learnt proposition E — is not an expert after all. You know that the E-person has come to basic-know E, and has updated as rationality requires, but you don't know that the E-person has come to *super*-know E. And this is because for all you know the E-person may not know that (s)he is the E-person, and so does not know that the proposition that (s)he has just learnt is E.

This explains why it is that whether you are rationally required to defer to an agent can depend on how that agent is designated (Bronfman, 2015; Mahtani, 2016). For example, you consider Alice an expert and so are required to defer to her because you know that she has come to know some proposition that you don't know, and that under any designator that you might use to pick this proposition out (e.g. 'the extra proposition that Alice has come to know'), Alice will recognize the proposition so designated. In contrast, you need not consider the E-person as an expert, and so are not required to defer to him or her, for even though you know there is a proposition that the E-person (basic-)knows and you don't, namely E, you don't expect the E-person to recognize the proposition so designated. More generally, you cannot designate any proposition that you don't know but that you do know the E-person has come to super-know. Thus you consider Alice an expert, but you don't consider the E-person an expert, even though Alice and the E-person may be one and the same.³²

Having shown how the distinction between basic-know and super-know can be used to formulate a plausible version of the Reflection Principle, I turn now to the Sure-Thing Principle.

6. The Sure-Thing Principle (Savage, 1954)

We can introduce the principle using as usual the World's Smallest Lottery as an example. Let's suppose that you are trying to decide whether to buy a cup of coffee, but you don't yet know whether you have won the lottery (i.e. whether WIN1 is true). To help you decide, you consider whether you would buy a cup of coffee if you were to come to know that WIN1 is true, and you conclude that you would; you then consider whether you would buy a cup of coffee if you were to come to know that WIN1 is not true, and you conclude that you would. According to the Sure-Thing Principle, then, you ought to

³² See (Bronfman, 2015; Mahtani, 2016) for further discussion of this point, and to see how this relates to the 'Stopping Time Restriction' on the Reflection Principle (Schervish, Seidenfeld, & Kadane, 2004).

just go ahead and decide to buy a cup of coffee without waiting to find out whether WIN1 is true.

Savage claimed that this sort of informal presentation of the principle involved ‘new undefined technical terms referring to knowledge and possibility’, which (in his view) rendered this informal presentation of the principle ‘mathematically useless without still more postulates governing these terms’ (Savage, 1954, p. 22). For this reason, Savage turned to a formal treatment, but there has nevertheless been plenty of interest in the informal presentation, including various objections (Jeffrey, 1982; McClennen, 1988). In this section I focus on a particular puzzle concerning the informally presented principle. To see this puzzle, first consider this ‘Generalized Sure-Thing Principle’:

If: for some action A, an agent knows that there is a proposition E such that if (s)he were to come to know E, then she would decide to do A

Then: that agent ought to decide to do A.

This principle seems like a sensible generalization of the reasoning behind the Sure-Thing Principle, and given that the Sure-Thing Principle seemed compelling, this more general principle may seem compelling too. However this principle is definitely false, as we can see once again by applying it in the case of the World’s Smallest Lottery. Suppose that — unlike in the above example — you would decide to buy a cup of coffee (perhaps to console yourself) if and only if your credence that you had won the lottery (i.e. your credence in WIN1) is strictly less than $\frac{1}{2}$. As things stand, you do not decide to buy a cup of coffee, because your credence in WIN1 is $\frac{1}{2}$. However, you know that there is a proposition — namely E (defined as usual as the proposition that is G) — such that if you were to come to know it, your credence in WIN1 would be $\frac{1}{3}$, and so you would decide to buy a cup of coffee. By our Generalized Sure-Thing Principle, then, you ought to decide at once to buy a cup of coffee. But that is clearly false, so the principle has led us astray here.³³

It has been noted that the Sure-Thing Principle is plausible only in cases where what might be learnt forms a partition (Aumann, Hart, & Perry, 2005). Thus we can locate the problem in the previous example by pointing out that the different propositions that you might come to know when you come to (basic-)know E do not form a partition. For example, E might be the proposition WIN2, or it might be the proposition WIN3, and these two propositions might both be true together. However this leaves us with the question of why it is that the principle is plausible only in cases where what might be learnt forms a partition. The Generalized Sure-Thing Principle looks at first sight compelling just as it is, without any restriction added. As Aumann et al. put it, considering a case where the possibilities do not form a partition, ‘on its face, the reasoning leading to the conclusion appears no less compelling than before’ (Aumann, Hart, & Perry, 2005, p. 8). How can we explain this?

The Generalized Sure-Thing Principle is compelling but false. As usual, the explanation for this turns on the distinction between basic-know and super-know. The Generalized Sure-Thing Principle is compelling because it is true under one reading — with the instance of ‘know’ immediately before ‘E’ read as super-know; yet it is false under another reading — with the instance of ‘know’ immediately before ‘E’ read as basic-

³³ Once again, we could show that the principle leads not just to counterintuitive conclusions, but to contradictory conclusions by extending the example.

know. We can see then why we were led astray when we applied the principle (for the second time) to the case of the World's Smallest Lottery. It is true that you know that if you were to come to basic-know E, then you would have a credence of $\frac{1}{3}$ in WIN1, and so would decide to buy a cup of coffee. But it is not true that you know that if you were to come to super-know E, then you would have a credence of $\frac{1}{3}$ in WIN1, and so would decide to buy a cup of coffee. For E might be WIN4, say, and then if you were to come to super-know E (i.e. to know that E is WIN4), you could at once infer WIN1, in which case you would not decide to buy a cup of coffee. Thus if we interpret the relevant instance of 'know' as super-know, the Generalized Sure-Thing Principle does not lead us astray.

I now finish by comparing my response to the Generalized Sure-Thing Principle to that of Aumann et al. (Aumann, Hart, & Perry, 2005). Aumann et al. point out that a rational agent pays attention not just to what evidence (s)he has gained, but also to *how* it has been gained. In their terminology, we can distinguish between an event E occurring, and an agent receiving some signal s that event E has occurred. Thus for example, in our case of the World's Smallest Lottery, we might distinguish between ticket 3's being a losing ticket (E), and the organizer telling you that ticket 3 is a losing ticket. Where an event E does not entail its signal s, you need to be sure to conditionalize not on E (the event's occurring), but rather on S (the fact that you have received the signal), taking into account whatever information you have about the signalling process.³⁴

Is this point enough to rescue the Generalized Sure-Thing Principle, without any need for my distinction between basic-know and super-know? I claim that it is not. As discussed in section 4, sometimes an agent needs to super-know a proposition in order to be able to bring to bear whatever evidence (s)he has about the relevant information-gathering process. Similarly, sometimes an agent needs to be able to identify a signal by name in order to bring to bear his or her information about the signalling function that produced it. We can see this in our counterexample to the Generalized Sure-Thing Principle in the case above. You know that if you were to come to (basic-)know just E, then you would have a credence of $\frac{1}{3}$ in WIN1, and so would decide to buy a cup of coffee. In this counterfactual situation you are updating on all your relevant evidence: you are taking into account any information you have (which may be none) about the relevant signalling function. Thus you are not lacking in rationality in this counterfactual situation: what you are lacking is knowledge that the proposition that is all that you have come to know — the signal that you have received — is E. Had you known that it was E, your credence in WIN1 would have adjusted accordingly, and the troublesome antecedent of the Generalized Sure-Thing Principle would have been false.

On my view then, the way to rescue the Generalized Sure-Thing Principle is to recognize that it is ambiguous as it stands. As Savage pointed out (Savage, 1954, p. 22) the principle involves terms such as 'knowledge' which are not fully defined. We have seen here how the ambiguity in the term 'knowledge' has left the principle open to two

³⁴ Aumann et al. provide a model to spell-out their position. They suggest that in cases where an agent receives a signal without knowing what the relevant signalling process was, the model is incoherent, and so presumably cannot provide any guidance (Aumann, Hart, & Perry, 2005, p. 8). However situations in which agents receive signals without being certain of the relevant signalling function are so common that it seems undesirable to exclude these situations from the account. I therefore assume a natural extension of the account of Aumann et al., and assume that a rational agent will take into account any evidence that (s)he has about the relevant signalling function — even if the agent is in a state of uncertainty about it.

different interpretations — one certainly false, and one possibly true. Here I have distinguished these two readings, and so unearthed a more plausible version of the Generalized Sure-Thing Principle.

Conclusion

The word ‘know’, when followed by the name for a proposition, can be interpreted in two ways: as basic-know or as super-know. Recognizing this distinction is the key to solving a range of puzzles, and clarifying several debates. Here I have shown how the distinction can be used to construct plausible versions of the Opaque Proposition Principle, the Reflection Principle, and the Sure-Thing Principle.

Bibliography

- Arntzenius, F. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy*, 100(7), 356–370.
- Aumann, R. J., Hart, S. and Perry, M. (2005). *Conditioning and the Sure-thing Principle*. Center for the Study of Rationality. Retrieved from <http://www.ma.huji.ac.il/raumann/pdf/dp393.pdf>
- Bovens, L. and Rabinowicz, W. (2011). Bets on Hats: On Dutch Books Against Groups, Degrees of Belief as Betting Rates, and Group-Reflection. *Episteme*, 8(3), 281–300.
- Bradley, D. (2012). Four Problems about Self-Locating Belief. *Philosophical Review*, 121(2), 149–177.
- Briggs, R. (2009). Distorted Reflection. 118(1), 59–85.
- Bronfman, A. (2015). Deference and Description. *Philosophical Studies*, 172, 1333–1353.
- Christensen, D. (1991). Clever Bookies and Coherent Beliefs. *Philosophical Review*, 100(2), 229–247.
- Elga, A. (2000). Self-locating belief and the Sleeping Beauty problem. *Analysis*, 60(2), 143–147.
- (2007). Reflection and Disagreement. *Nous*, 41(3), 478–502.
- Evans, G. (1979). Reference and Contingency. *The Monist*, 62(2), 161–189.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Anthony, *Philosophers Without God: Meditations on Atheism and the Secular Life* (pp. 194–214). Oxford: OUP.
- Fitelson, B. (2012). Evidence of evidence is not (necessarily) evidence. *Analysis*, 72(1), 85–88.
- Gaifman, H. (1988). A Theory of Higher Order Probabilities. In B. Skyrms and W. L. Harper, *Causation, Chance and Credence: Proceedings of the Irvine Conference on Probability and Causation Volume 1* (pp. 191–219). Dordrecht: Springer Netherlands.
- Hájek, A. (2005). The Cable Guy Paradox. *Analysis*, 65(286), 112–119.
- Hall, N. (1994). Correcting the Guide to Objective Chance. *Mind*, 103, 505–18.
- Hedden, B. (2015). Time-Slice Rationality. *Mind*, 124(494), 449–491.
- Hutchison, K. (1999). What Are Conditional Probabilities Conditional Upon? *British Journal for the Philosophy of Science*, 50(4), 665–695.
- Jeffrey, R. (1982). The Sure Thing Principle. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 2, 719–739.

- King, J., Soames, S. and Speaks, J. (2014). *New Thinking about Propositions*. Oxford: OUP.
- Korb, K. B. (1994). Infinitely many resolutions of Hempel's paradox. *Proceedings of the 5th conference on Theoretical aspects of reasoning about knowledge* (pp. 138–149). San Francisco: Morgan Kaufmann Publishers Inc.
- Kotzen, M. (2012). Selection Biases in Likelihood Arguments. *British Journal for the Philosophy of Science*, 1–15.
- Mahtani, A. (2014, March 22). Deference and Designators. Retrieved from <http://www.sas.ac.uk/videos-and-podcasts/philosophy/deference-and-designators>
- (2016). Deference, Respect and Intensionality. *Philosophical Studies*.
- McClennen, E. (1988). Sure-thing doubts. In P. Gardenfors and N.-E. Sahlin, *Decision, Probability and Utility* (pp. 166–182). Cambridge: CUP.
- Roche, W. (2014). Evidence of evidence is evidence under screening-off. *Episteme*, 11 (1), 119–124.
- Salmon, N. (1989). Illogical Belief. *Philosophical Perspectives*, 3, 243–285.
- Savage, L. (1954). *The Foundations of Statistics*. New York: Wiley.
- Savant, M. v. (1990). Ask Marilyn. *Parade Magazine*, p. 16.
- Schervish, M. J., Seidenfeld, T. and Kadane, J. B. (2004). Stopping to Reflect. *Journal of Philosophy*, 101(6), 315–322.
- Selvin, S. (1975). A Problem in Probability (Letter to the Editor). *The American Statistician*, 29(1), 67.
- Tal, E. and Comesaña, J. (2015). Is Evidence of Evidence Evidence? *Noûs*, 1–18.
- Talbott, W. J. (1991). Two Principles of Bayesian Epistemology. *Philosophical Studies*, 62, 135–150.
- Titelbaum, M. (2013). *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford: OUP.
- van Fraassen, B. (1984). Belief and the Will. *The Journal of Philosophy*, 81(5), 235–256.
- Weisberg, J. (2005). Conditionalization, Reflection, and Self Knowledge. *Philosophical Studies*, 135, 179–197.