

Transparency and Partial Beliefs

Weng Hong Tang

National University of Singapore

1 Introduction

We rely on perception to give us knowledge about our environment. We look out the window to see if it's raining, prick our ears to determine if our pursuers are approaching, or sniff the milk to find out if it has gone bad. But we know various things, not only about our environment, but about our inner lives as well. We may know, for instance, that we *believe* it's raining, that we *fear* our pursuers will find us, and that we *desire* not to drink milk that's gone bad.

We also seem to have a special kind of access to our inner lives that we lack with respect to other people's inner lives. For instance, as Byrne (2011) puts it, such access is *peculiar*—‘one has a way of knowing about one's mental states that one cannot use to know about the mental states of others’ (p. 202). Intuitively, whereas you can find out through introspection whether you believe it'll rain, you can't find out in the same manner whether I believe it'll rain.¹

How do we gain such self-knowledge? Armstrong (1968) endorses Kant's view that self-knowledge is due to ‘the operation of “inner sense”’ (p. 95). Further, he thinks that this operation amounts to ‘perception of the mental’—just as we have mental faculties for perceiving our environment, we've a mental faculty for perceiving mental states or processes (Armstrong 1981, p. 60). But as Byrne (2005) observes, such a view ‘is not infrequently taken to be a crass mistake’ (p. 80). Moran

¹This is not to say that we can't learn about our inner lives in the same way that we learn about the inner lives of others—for instance, through observing the relevant behaviour or through testimony.

(2001), for example, speaks of the ‘original embarrassment of the “inner eye” and the concern that it cannot be cashed out as anything other than a misleading metaphor’ (p. 13).

According to Evans (1982), we don’t need to appeal to an inner sense to find out what we believe. He writes:

[I]n making a self-ascription of a belief, one’s eyes are ... directed outward—upon the world. If someone asks me ‘Do you think there is going to be a third world war?’, I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question ‘Will there be a third world war?’ I get myself in a position to answer the question whether I believe that p by putting into operation whatever procedure I have for answering the question whether p . (Evans 1982, p. 225)

If Evans is right, we can explain why our access to our own mental states is peculiar. While answering ‘yes’ to the question ‘Will there be a third world war?’ may tell me that I believe there’ll be a third world war, it won’t typically tell me whether someone else has such a belief (Byrne 2011, p. 207).²

Now, the claim that we *must* answer the question ‘Do you believe p ?’ by answering the question ‘Is it the case that p ?’ is too strong. To use Gertler’s (2011) example, suppose I’m asked whether I believe there’s going to be a third world war. Recalling that I told a friend this morning that I feared such a war was imminent, I answer ‘yes’. In doing so, I’ve answered the question ‘Do you believe there’ll be a third world war?’, but not via answering the question ‘Will there be a third world

²Byrne (2011) takes the upshot of Evans’s remarks to be that we obtain self-knowledge by reasoning in accord to the ‘doxastic schema’, i.e., by inferring that we believe p from the premise that p (p. 204). He also thinks that if Evans is right, we can explain why our access to our own mental states is *privileged*—we can explain why, ‘by and large, beliefs about one’s mental states are more likely to amount to knowledge than one’s corresponding beliefs about others’ mental states’ (*ibid.*, p. 202). Byrne argues for this point in part by holding that when one ‘infers that one believes that p from the premiss that p , then one’s second-order belief [i.e., one’s belief that one believes that p] is true, because inference from a premiss entails belief in that premiss’ (*ibid.*, p. 206). He also thinks that, typically, such a second-order belief is *safe* (i.e., it could not easily have been false). For more details, see Byrne (2005) (pp. 96–8).

war?’

But presumably, I *can* answer the former question by answering the latter. Answering ‘yes’ to the latter will presumably tell me that I believe there’ll be a third world war, whereas answering ‘no’ or something along the lines of ‘It can go either way’ will presumably tell me that I don’t. Generally, one may be sympathetic to the weaker claim that we *can* find out whether we believe p simply by asking ourselves whether p . Henceforth, I’ll focus on this weaker claim—call it *transparency*.

In particular, I’ll focus on transparency with respect to *partial belief*. So far, I’ve been treating belief as an all-or-nothing attitude: either we believe a proposition or we don’t, in which case either we disbelieve it or we suspend judgement about it. But arguably, belief comes in degrees too. For instance, we may divide our confidence equally between a coin landing heads and not landing heads; we may be three times as confident that we’ll roll an even number as we are that we’ll roll a ‘5’; we may be 80% confident that it’ll rain tomorrow and 10% confident that there’ll be a thunderstorm. Given that we have self-knowledge of our partial beliefs—that, for example, we can find out introspectively that we are quite confident, though not absolutely certain, that it’ll rain tomorrow—a natural question arises: are partial beliefs transparent?³

In this paper, I clarify the question above. I also consider various attempts to answer the question in the affirmative. To anticipate, my verdict is pessimistic: I argue that such attempts fail. This failure puts pressure on those who claim that all-or-nothing beliefs are transparent and that their transparency helps account for our self-knowledge of them.

Why would there be such pressure? Some philosophers, for example, Byrne

³Some philosophers, e.g., Harman (1986) and Pollock (2006), hold that talk of partial beliefs can be replaced by talk of all-or-nothing beliefs (p. 24; p. 94). For example, on such a view, a degree of confidence of 0.8 in p can be thought of as an all-or-nothing belief that the objective probability of p ’s being true is 0.8. If the view is correct, a question about the transparency of partial beliefs really boils down to a question about the transparency of a subset of all-or-nothing beliefs. But it’s worth pursuing the question whether partial beliefs are transparent on the assumption that such a view is false. After all, several philosophers find the view implausible, and there are compelling arguments against it—see, for example, Christensen (2004) and Frankish (2009) (pp. 18–20; pp. 77–8).

(2011), have explored the possibility that other propositional attitudes such as desire and intention are transparent too. We may wonder, for instance, whether a question about what we desire may be answered via a question about what is objectively desirable. Though it's possible that beliefs are transparent whereas attitudes like intentions and desires are not, Byrne (2005) thinks that such a situation would raise a puzzle (p. 99). Suppose you think that the transparency of beliefs helps us account for self-knowledge without our having to posit an inner sense. But suppose that desires and intentions—of which we plausibly possess self-knowledge—are not transparent. Then we might still need to appeal to an inner sense to account for such self-knowledge, and one may wonder why an inner sense shouldn't be operating in the case of beliefs (*ibid.*).

A similar point can be made if all-or-nothing beliefs are transparent but partial beliefs are not. After all, we do possess self-knowledge of partial beliefs—we often know whether we are very confident, or slightly confident, or rather doubtful that a proposition is true. Perhaps we might not often know our exact level of confidence in a proposition. But even if, for example, I do not know that I'm exactly 50% confident that a coin will land heads, I may know that I'm roughly 50% confident that it will do so. So if partial beliefs are not transparent, we might still need to appeal to an inner sense to account for our self-knowledge of such mental states, and one may wonder why an inner sense shouldn't be operating in the case of all-or-nothing beliefs.

Granted, although Evans (1982), Byrne (2011), and Moran (2001) think that transparency lets us account for self-knowledge without our having to posit an inner sense, there are accounts of self-knowledge that neither appeal to an inner sense nor to transparency.⁴ So it is by no means clear that we have to appeal to an inner sense to account for self-knowledge of attitudes like desires and intentions

⁴For instance, according to Paul (2012), 'one can come to know what one intends by self-ascribing the content of a decision about what to do as one's intended action' (p. 335). In accounting for self-knowledge of our intentions, Paul does not appeal to an inner sense. Neither does she appeal to transparency; she thinks that 'intentions just aren't transparent enough' (*ibid.*, p. 329). See Gertler (2015) for a discussion and an overview of several other accounts of self-knowledge.

if it turns out that they are not transparent. Further, there is no inconsistency in holding that self-knowledge of beliefs is explained by transparency whereas self-knowledge of desires and intentions is explained by some other account. Perhaps beliefs are different enough from desires and intentions that it is not too surprising if self-knowledge of the latter requires different treatment from self-knowledge of the former. But given that partial beliefs and all-or-nothing beliefs are both doxastic attitudes, a failure of transparency with respect to partial beliefs should prove more worrying than a failure of transparency with respect to desires or intentions. At least, it should put some pressure on those who think that all-or-nothing beliefs are transparent but that partial beliefs are not to explain why the two kinds of belief should differ with respect to transparency.

Given the preceding, it's somewhat surprising that the issue of whether partial beliefs are transparent has received no or little attention in the literature on transparency and self-knowledge. This paper helps fill the gap; the issue deserves greater attention.

2 The question of transparency for partial beliefs

What are we asking when we ask whether partial beliefs are transparent? Before addressing this question, let's consider a qualification that one might wish to add to the claim that all-or-nothing beliefs are transparent. Byrne (2005) writes that 'the correct formulation of Evans's observation is that one can *typically* answer the question "Do I believe P?" simply by considering whether P' (p. 82; Byrne's emphasis). For he agrees with Moran (2001) that there are *atypical* cases in which one's beliefs are not transparent.⁵

To use Moran's example, suppose a patient undergoing therapy has repressed her belief that she's been betrayed by a friend. In response to the question 'Did your friend betray you?', she'll answer 'no' or say that she can't settle the question either way. Yet, through therapy, even while the belief remains repressed, she might

⁵For my purposes, I'll grant the qualification that Moran and Byrne make.

become convinced, ‘through the eliciting and interpretation of evidence of various kinds’, that she has such a belief (*ibid.*, p. 85). In such a case, she will answer ‘yes’ to the question ‘Do you believe that your friend betrayed you?’

Following Byrne (2005), we may say that the patient’s repressed belief is *alienated*, where a belief is alienated if and only if it is ‘to a significant extent inferentially isolated’ and ‘not expressible by the subject in unembedded speech’ (p. 87). If Moran and Byrne are right, alienated beliefs are not transparent. But cases involving alienated beliefs are atypical and are cases in which the normal capacity for self-knowledge is impaired. So they pose no problem for the claim that beliefs are *typically* transparent, and they pose no problem for the accompanying claim that such transparency helps account for self-knowledge.

Are there other atypical cases in which a failure of transparency poses no problem for such claims? What about unalienated beliefs that fall short of ideal rationality? I take it that the thesis of transparency is not meant to apply to ideally rational agents only. After all, the thesis is supposed to help account for *our* capacity for self-knowledge, and we are less than ideal. So we shouldn’t count a case as atypical just because it involves a less than ideally rational belief. Now, it’s beyond the scope of this paper to provide an exhaustive taxonomy of typical and atypical cases. But for my purposes, given the work that the thesis of transparency is meant to do, a case should be dismissed as atypical only if it involves alienated or irrational beliefs *that impair our normal capacity for self-knowledge*.⁶ If there’s a case in which our capacity for self-knowledge is functioning normally but in which transparency fails, it won’t do to dismiss the case as atypical. Instead, the case should be taken to provide evidence that self-knowledge is not best explained by transparency.

If all-or-nothing belief is transparent, we can typically answer the question

⁶And perhaps, even some cases that involve downright irrational beliefs, for example, beliefs that are inadequately supported by our evidence, shouldn’t count as atypical. On the face of it, a belief of ours that is inadequately supported by the evidence (but unalienated) is no bar to our having self-knowledge of it. The same scant evidence that leads to our believing *p* might also lead to our answering ‘yes’ to the question ‘Is it the case that *p*?’

‘Do I believe p ?’ by answering the question ‘Is it the case that p ?’ In doing so, we are in the first instance answering a question about the content of the relevant belief, and not about the belief itself—we are in the first instance concerned about whether p and not about whether we *believe* p . But now suppose we wish to find out how confident we are that p . If partial belief is transparent, we should typically be able to answer questions such as ‘How confident am I that p ?’ or ‘Am I 80% confident that p ?’ via answering other questions. As in the case of all-or-nothing belief, these other questions should in the first instance be about the content of the relevant partial belief, and not about the partial belief itself.

What might these questions be? To find out how confident we are that p , may we ask whether p ? The suggestion won’t do if we restrict ourselves to ‘yes’ or ‘no’ answers. For answering ‘yes’ or ‘no’ to the question ‘Is it the case that p ?’ may tell us that we believe p or that we don’t, but it won’t tell us how confident we are that p . To deal with this problem, one might suggest we admit probabilistic answers such as ‘It’s 80% likely that p ’ or ‘The probability that p is 0.6’. According to Evans (1982), in self-ascribing all-or-nothing belief, ‘the subject’s concentration ... is on the ... world: how does he, or would he, judge it to be’ (p. 230)? One might think that, in self-ascribing partial belief, instead of judging whether the world is a certain way or not, we judge how *likely* it is to be a certain way.

Admitting probabilistic answers is a move in the right direction. But such answers are not, strictly speaking, answers to the question ‘Is it the case that p ?’ Suppose you respond to it by simply saying, ‘The probability of p is 0.99’. One may protest that you’re not answering the question and say, ‘Well, you’ve told me that p is very *probable*—but what I want to know is whether it’s *true*’. There is, however, a simple fix to this worry: instead of asking ‘Is it the case that p ?’ ask ‘How likely is it that p ?’ or ‘What is the probability of p ?’ If partial belief is transparent, an answer of the form ‘It’s $x\%$ likely that p ’ should tell us that we are $x\%$ confident that p .⁷

⁷As mentioned earlier, one might think that we often do not know exactly how confident we are that something is true. For instance, even when we are exactly 99.85% confident that p , we might

But suppose that in response to the question ‘How likely is it that p ?’, I answer ‘It’s $x\%$ likely that p ’. And suppose ‘It’s $x\%$ likely that p ’ means the same as ‘I’m $x\%$ confident that p ’. Then even if my answer tells me correctly how confident I am that p , its doing so provides no support for the view that partial belief is transparent. For in giving such an answer, I would be reporting on my state of confidence and not on the content of my partial belief. In such a case, answering the question ‘How likely is it that p ?’ amounts to answering the question ‘How confident are you that p ?’ A similar problem arises if we suppose that expressivism with respect to utterances of the form ‘It’s $x\%$ likely that p ’ is correct.⁸ Again, even if my answer tells me correctly how confident I am that p , no support accrues to the view that partial belief is transparent. For in such a case, I am not reporting on or making a judgement about p ; I’m merely expressing my state of confidence in p .

The above shows that in stating what it is for partial belief to be transparent, we should take ‘likelihood’ or ‘probability’ to refer to some kind of objective probability. Then a claim of the form ‘It’s $x\%$ likely that p ’ will be about p instead of about one’s state of confidence. It remains to be seen, however, whether partial belief is transparent.

3 Is Partial Belief Transparent?

If partial belief is transparent, we may answer the question ‘How confident are you that p ?’ by answering the question ‘How likely is it that p ?’ (or ‘What’s the probability of p ?’), where ‘likely’ should be taken to mean *objectively likely*. But since there are different ways to understand what it is for something to be objectively likely, there are different ways to understand what it is for partial belief

not know that we are so; perhaps, at best, we know that we are roughly 99% confident that p . This shouldn’t pose a problem for the view that the transparency of partial beliefs helps account for our self-knowledge of them. One may hold that judging that it’s $x\%$ or roughly $x\%$ likely that p should tell us that we are roughly $x\%$ confident that p . Or, speaking in qualitative terms, one may hold that judging that it’s quite probable that p should tell us that we’re quite confident that p . For convenience, I’ll continue to speak as if we often know exactly how confident we are of various propositions being true. Nothing of importance hinges on this simplifying assumption.

⁸See Yalcin (2007) for more on expressivism with respect to credences.

to be transparent. In what follows, I consider some of these ways. For brevity, let ‘transparency’ refer to the transparency of partial belief unless otherwise stated.

3.1 Transparency: an Appeal to Relative Frequencies

Let’s say that the relevant notion of probability is relative frequency. Then, given transparency, we should be able to answer the question ‘How confident are you that the coin will land heads?’ by answering the question ‘What’s the relative frequency of the coin’s landing heads?’ For example, if our answer to the latter question is 0.8, then transparency should tell us that our confidence in the coin’s landing heads is 0.8 too.

This view faces at least two problems. First, suppose Ida is 80% confident that the universe originated with the Big Bang, and suppose she knows the preceding to be true of herself. Since the universe’s coming into existence is a one-off, non-repeatable event, Ida might think that it makes no sense to talk about its relative frequency. In such a case, she will think that the question ‘What is the relative frequency of the universe originating with the Big Bang?’ is ill-posed. Or she may judge that the relevant relative frequency is either 1 or 0 insofar as she thinks that the relative frequency of a one-off event having happened is either 1 or 0, depending on whether it happened or not. In either case, there’s a problem for transparency. Ida needn’t have any alienated nor irrational beliefs that impair her capacity for self-knowledge. But her answer to the preceding question won’t tell her how confident she is that the universe originated with the Big Bang.

Second, although we sometimes base our confidence in a proposition on what we judge the corresponding relative frequency to be, we are often unable to do so. For we are often uncertain—and rationally so—about the relative frequency of an event happening. Such uncertainty then results in a mismatch between our confidence in the event happening and what we judge the corresponding relative frequency to be. For example, suppose Ida knows that a coin—let’s call it C —is either two-headed or two-tailed but does not know which. Suppose her degree of

confidence in C landing heads is 0.5, as seems rational. But she may think that the relative frequency of C landing heads is either 1 or 0, depending on whether C is two-headed or two-tailed. In such a case, her answer to the question ‘What’s the relative frequency of the coin’s landing heads?’ won’t yield an answer to the question ‘How confident are you that the coin will land heads?’ Transparency fails, but the case shouldn’t be dismissed as atypical, since Ida needn’t have any alienated or irrational beliefs that impair her capacity for self-knowledge.

3.2 Transparency: An Appeal to Objective Chance or Propensity

To avoid the first problem above, we may take ‘probability’ to refer to objective chance or propensity instead of relative frequency. For it makes sense to talk about the chance or propensity of a one-off event happening. For instance, on the current proposal, the probability of the universe originating with the Big Bang may have an intermediate value.

But although this proposal avoids the first problem, it succumbs to a version of the second problem. Although we sometimes base our confidence in a proposition on what we judge the corresponding chance or propensity to be, we are often unable to do so. For we are often uncertain—and rationally so—about the chance or propensity of a proposition’s being true. Such uncertainty then results in a mismatch between our confidence in the proposition and what we judge the corresponding chance or propensity to be. For example, suppose Ida knows that C is either two-headed or two-tailed, but she does not know which. Suppose her degree of confidence in C landing heads is 0.5, as seems rational. But she may think that the objective chance or propensity of C ’s landing heads is either 1 or 0, depending on whether C is two-headed or two-tailed. In such a case, she can’t answer the question ‘How confident are you that the coin will land heads?’ by answering the question ‘What’s the objective chance or propensity of the coin’s landing heads?’ Again, transparency fails. And again, the case shouldn’t be dismissed as atypical—Ida needn’t have any alienated or irrational beliefs that impair her capacity for

self-knowledge.

For good measure, here's another example. Suppose Jack has no clue as to whether p and no clue as to the objective chance (or propensity) of p being true. (The example will also work if the notion of probability invoked is relative frequency.) Along comes Sue, whom Jack knows is 60% reliable when it comes to making true statements. Sue asserts that p and nothing else. It seems reasonable for Jack to become 60% confident that p —let's suppose that he does. But if he's asked, 'What's the objective chance of p ?', it also seems reasonable for him to say something along the lines of 'I don't know'. In fact, he shouldn't answer '0.6'. This is because Sue's being 60% reliable does not mean that any proposition she asserts has a 0.6 objective chance of being true. For example, Sue's being 60% reliable is compatible with a case in which, 60% of the time, she asserts propositions whose objective chance of being true is 1, and 40% of the time, she asserts propositions whose objective chance of being true is 0.⁹

⁹One might think that the proposal faces another worry. Suppose a coin was tossed yesterday. Knowing that it's a fair coin but knowing nothing else relevant about the outcome of the toss, Ida is 50% confident that the coin landed heads. On the current proposal, she may find out that she is so by answering the question 'What's the objective chance or propensity of the coin's having landed heads?' But one might think that she can't if we suppose that, like philosophers such as Lewis (1980) and Schaffer (2007), she thinks that the objective chance of a past event having happened is either 1 or 0, depending on whether it happened or not. Similarly, she may think that the propensity of the coin's having landed heads yesterday is either 1 (if it landed heads) or 0 (if it did not). In fact, she may even reject the question 'What is the propensity of the coin's having landed heads?' as ill-posed if she thinks that it makes no sense to talk about the propensities of past events having happened. In general, the worry is that we often assign intermediate degrees of confidence to various past events having happened—and are rational in doing so. But if asked about the chance or propensity of those events having happened, we may be inclined to answer either '1' or '0'—provided we don't think the question is ill-posed.

This is a very natural worry to raise, but proponents of transparency might ultimately be able to answer the worry. To begin, they may grant that having certain theoretical commitments to do with the nature of probability might corrupt the transparency method, which lets us find out how confident we are that p by asking ourselves how likely it is that p . But they might suggest that most people—the folk—don't usually have such sophisticated commitments, and are happy to assign intermediate chances to past events. That someone like Ida—or Lewis—might have to rely on some other means to obtain self-knowledge does not undermine the claim that transparency offers the primary explanation of how the folk obtain self-knowledge. Or to use an analogy offered by an anonymous referee, it is not a strike against an account of how the folk generally attribute moral responsibility if the account can't explain the reasoning of fatalists who think that, ultimately, none of us are morally responsible for our actions. To push the analogy further, it might be that, as Stace (1952) observes, even philosophers who avow that fatalism is true or that there's no such thing as free will can't help but reflexively attribute moral responsibility in everyday situations (pp. 248–9). Similarly, even though Ida and Lewis might have certain theoretical commitments regarding the nature of probability, such commitments might not affect their using the transparency method to acquire self-knowledge in everyday situations. When asked, 'How confident are you that it rained

3.3 Transparency: an Appeal to Epistemic Probability

What if, instead of appealing to chance, propensity, or relative frequency, we appeal to epistemic (or logical) probability? We can think of the epistemic probability of p as the degree of support that our total relevant evidence e confers upon p . For example, the epistemic probability that all emeralds are green given that all observed emeralds are green is high; the epistemic probability that the next apple we pick will be rotten given that, so far, half of the twenty apples we've randomly picked are rotten is of some middling value; and in the extreme, the epistemic probability that a logical or mathematical truth is true (conditional on any proposition whose probability is greater than 0) is 1, whereas the epistemic probability that a logical or mathematical falsehood is true (conditional on any proposition whose probability is greater than 0) is 0.

Appealing to epistemic probabilities helps us avoid the above worries faced by an appeal to chance, propensity, or relative frequency. First, the epistemic probability of a one-off, non-repeatable event may be of intermediate value. Ida may judge that the epistemic probability of the Big Bang theory being true given her total relevant evidence is 0.8. In fact, her judging so may explain her being 80% confident that the theory is true. Second, suppose Ida knows that C is either two-headed or two-tailed, but does not know which. It seems that the epistemic

yesterday?', Ida might, at a non-reflective level, answer the question via the question 'What's the objective chance that it rained yesterday?', relying *tacitly* on a pre-theoretical notion of objective chance on which past events may be assigned intermediate probabilities. To be clear, this is not to deny that if Ida considers the latter question explicitly, her answer to it will be either '0' or '1'. But the proponent of transparency can maintain that even if Ida is unwilling or unable to use the transparency method explicitly, she nonetheless gains self-knowledge of her partial beliefs by her tacit, non-reflective use of the method. She needn't be aware that she's using the method, just as one may speak grammatical English without being aware of whether one is following the rules of English grammar correctly—cf. Davies (2015) on Chomsky and tacit knowledge. (Thanks to an anonymous referee for suggesting the above response on behalf of the proponent of transparency.)

Whether the response above is ultimately successful depends in part on whether the folk indeed think (tacitly if not explicitly) that past events can be assigned intermediate objective chances. If the folk think that the chance of a past event should be either 0 or 1 in order to reflect the thought that the past is in some sense closed—if the folk notion of chance is akin to Lewis's—then the response won't work. But I'm not sure what the folk really think, and the project of investigating the folk notion of objective chance is too large for me to carry out here. This being so, I grant the proponent of transparency that, as it stands, the worry raised at the beginning of this footnote is not entirely compelling.

probability of C 's landing heads given her relevant evidence should equal 0.5. In fact, her judging so may explain her being 50% confident that C will land heads. Third, suppose Jack has no idea as to whether p , and Sue, whom Jack knows to be 60% reliable, asserts that p . It seems that the epistemic probability of p given Jack's relevant evidence should equal 0.6. In fact, his judging so may explain his being 60% confident that p upon hearing Sue assert p .

In general, on the current proposal, transparency says that we may answer the question 'How confident are you that p ?' by answering the question 'What is the epistemic probability of p given all your relevant evidence?' But though the proposal solves the worries above, and therefore seems promising, it faces some worries of its own.

Suppose we wish to answer the question 'How confident are you that it'll rain?' On the current proposal, we may do this by answering the question 'What is the epistemic probability of rain given all your relevant evidence?' But one might worry that we count something as being our evidence only if we believe it to be true, and so, to answer the question, we need to enquire as to what we believe.¹⁰

This is a natural worry, but perhaps one might try to resist it as follows. Suppose we are enquiring as to whether the streets are wet, and we reason as follows:

It's raining. And if it's raining, the streets are wet. So the streets are wet.

Let's grant that our reasoning to the conclusion that the streets are wet relies on our prior belief that it's raining and our prior belief that the streets are wet if it's raining—let's grant that we may only invoke premises that we believe. But note that the premises invoked are not about our internal doxastic states. For instance, there's no appeal to any premise of the form 'I believe that such-and-such'. Further, if we are asked, 'What evidence or reasons do you have for claiming that the streets are wet?', we may answer, 'Well, for one thing, it's raining'. And it seems that we can give such an answer without first asking ourselves whether we believe it's

¹⁰Cf. Gertler (2011). See, in particular, pp. 132–3.

raining.

Now, suppose we are asked ‘How confident are you that it’ll rain?’, and our relevant evidence for rain has to do with the sky being dark. According to the proposal under consideration, we can answer the question by answering the following one: ‘What is the epistemic probability of rain given that the sky is dark?’ Granted, to ask the right question, we first need to ask ourselves, ‘What is the relevant evidence?’ But analogously to the reasoning example, this is consistent with our enquiry not being about our internal doxastic states. First, one might argue that when enquiring about the epistemic probability of rain given the relevant evidence, we are not enquiring about our doxastic states per se; instead we’re asking about the degree of support that one proposition confers upon another. Second, one might argue that in response to the question, ‘What’s the relevant evidence?’, we may simply answer, ‘Well, the fact that the sky is dark’; and we may do this without first asking ourselves whether we believe that the sky is dark.

The response might be correct as far as it goes—perhaps answering a question about epistemic probability does not necessarily require us to answer questions that are directly about our mental states.¹¹ But there’s a related worry that the response cannot handle. Often we believe a proposition and are justified in believing it even if we can no longer recall our original source of justification for it and even if we have not gained any new evidence for believing the proposition.¹² This is true not only with respect to all-or-nothing belief but with respect to partial belief as well. Suppose that some time ago I acquired an all-or-nothing belief that q when a reliable person whom I trusted told me that q . Suppose that my memory has retained this all-or-nothing belief all this while. My believing that q (and my being

¹¹Perhaps we may resist the response if we hold that evidence is cashed out in terms of mental states. (See, for example, Feldman and Conee 2001, p. 2.) Consider, for instance, a view on which one’s evidence is just one’s beliefs. Then to ask ‘What’s the relevant evidence?’ is just to ask ‘What are our relevant beliefs?’ Consequently, if finding out how confident we are that something is true involves asking ourselves what our evidence is—if it involves asking ourselves what our *beliefs* are—then transparency is violated.

¹²Several epistemologists have made this point. For instance, Goldman (1999) writes that ‘[m]any justified beliefs are ones for which an agent once had adequate evidence that she subsequently forgot’ (p. 280). Also, see Senor 1993, p. 454; Audi 1995, p. 33; and Bernecker 2008, p. 114.

justified in believing that q) is compatible with my not remembering how I came to have the belief. It's compatible with my not being able to back up the belief with my original evidence for it and compatible with my not acquiring any new evidence for the belief. Similarly, suppose that some time ago I was 50% confident that p but then acquired a confidence of 80% in p when a person whom I knew to be fairly reliable told me that p . Today, although my memories may have faded a little, I'm still more than 50% confident that p —let's suppose that I'm 70% confident that p . This is compatible with my having forgotten my original evidence for p and compatible with my not having gained any new evidence for p .

But now, suppose I try to answer the question 'How confident am I that p ?' by answering the question 'What is the epistemic probability of p given all my relevant evidence?' I won't be able to do so if I've forgotten my original evidence for p and haven't gained any other evidence for it. The answer to the first question won't match the answer to the second.

Perhaps one might suggest that even if my original evidence for p is no longer available, I still have access to some other kind of evidence. In particular, one might suggest that my partial belief that p is supported by how clearly and distinctly my memory seems to indicate to me that p . For example, one might suggest that if I have a very clear and distinct memory of p being true, then that shows that I'm very confident that p , whereas if I have a somewhat less clear and distinct memory of p being true, then that shows that I'm somewhat less confident that p .¹³ Now, how clear and distinct my memory is may indeed indicate to me how confident I am of something being true. But the suggestion goes against the spirit of transparency—if I need to introspect and enquire about my own memory in order to answer a question about how confident I am that p , then I am not able

¹³If having a memorial belief that p involves having certain mental images, then one might think that how clearly and distinctly our memory indicates to us that p is a matter of how vivid or forceful such images are. But Alston (2005) thinks that even when one's memorial belief that p isn't based on any mental images, it could still be based on 'a sense of "pastness", a sense that what [one is] believing to have occurred is an experience [one] really had in the past' (p. 88). Following Alston's lead, one might suggest that, all things being equal, how strongly this sense of pastness is felt will indicate to one how confident one is that p .

to answer the question by focusing my attention solely on the question of whether p or on the question of how objectively likely it is that p .¹⁴

A remaining worry about the appeal to epistemic probability stems from our being less than absolutely confident as to what the relevant epistemic probabilities are. Earlier, we've seen how an appeal to epistemic probability might help deal with the case in which we know that a coin is either two-headed or two-tailed without knowing which. In such a case, we may be 50% confident that the coin will land heads and think that the epistemic probability of the coin landing heads given the relevant evidence is 0.5—even while judging that the chance, propensity or relative frequency of the coin's landing heads is either 0 or 1. But there are other cases in which our degrees of confidence do not match what we judge the relevant epistemic probabilities to be, because of our uncertainty regarding the latter.

Not being ideally rational, Ida is somewhat confident, though not absolutely sure, that S , a certain complex mathematical statement, is true. And she knows this about herself. But given that mathematical statements are either necessarily true or necessarily false, Ida also judges that the epistemic probability of the statement's being true is either 1 or 0—she just doesn't know which. In such a case, her answering the question 'What's the epistemic probability of S 's being true?' will fail to yield an answer to the question 'How confident are you that S ?'—her partial belief about S is not transparent even though her capacity for self-knowledge is intact.

A similar point can be made with respect to contingent propositions. Being non-ideal, we lack certainty as to the epistemic probabilities of various mathematical statements being true. But being non-ideal, we also lack certainty as to the epistemic probabilities of various contingent propositions being true. Callum, a trainee

¹⁴So long as we take epistemic probability to be a kind of probability that is conditional upon our evidence, the worry above arises whether we understand epistemic probability along Carnapian lines (see Carnap 1950), whether we take the notion of epistemic probability to be primitive (see Maher 2006), or whether we understand epistemic probability in terms of rational degrees of confidence. For instance, the argument above will work even if we substitute 'What degree of confidence is it rational for me to have in p given all my relevant evidence?' for 'What is the epistemic probability of p given all my relevant evidence?'

doctor, has available to him a body of evidence regarding whether his patient's parotid tumour is benign or malignant. Being an inexperienced trainee, Callum isn't sure what the evidence indicates.¹⁵ So he assigns an intermediate degree of confidence, say 0.5, to the patient's tumour being benign. Siti, a senior doctor, has available to her the same body of evidence. She declares that the patient's tumour is very likely to be benign and then walks away quickly to attend to an emergency. Unfortunately for Callum, due to a loud noise in the background, he isn't sure whether Siti, whom he trusts, said 'The patient's tumour is very likely to be benign' or 'The patient's tumour is very *unlikely* to be benign'. Consequently, even though he's now quite sure that the epistemic probability of the tumour being benign given the available evidence is either very high or very low, his confidence in the tumour being benign remains intermediate in value—and he knows this.¹⁶ In such a case, his answering the question 'What's the epistemic probability of the patient's tumour being benign?' will fail to yield an answer to the question 'How confident are you that the patient's tumour is benign?'—the partial belief in question is not transparent even though Callum's capacity for self-knowledge remains intact.¹⁷

¹⁵To clarify, I'm assuming that a piece of evidence may be *relevant* even if we do not grasp its evidential force. But denying this assumption won't help the proponent of transparency. Suppose we want to find out how confident we are that it will rain by asking 'What is the epistemic probability of rain given our total relevant evidence?' As mentioned earlier, we may arguably answer the question 'What's the relevant evidence?' without enquiring about our own mental states—arguably, the latter question can be thought of as a question about the external world. But suppose that we cash out *relevance* in a way such that to say that a piece of evidence is relevant is to say, in part, that we *grasp* its evidential force. Then the question 'What's the relevant evidence?' is in part a question about our mental states. Transparency is violated if we need to answer this latter question in order to find out how confident we are that it will rain. (Thanks to an anonymous referee for suggesting that I make this clarification.)

¹⁶The information that Siti either thinks the tumour is very likely to be benign or thinks it's very unlikely to be benign doesn't add much—if anything at all—to the original body of evidence. In particular, if the original body of evidence confers a high (or low) degree of support upon the proposition that the tumour is benign, then the original body of evidence plus this new piece of evidence should still confer a high (or low) degree of support upon the proposition (and not a degree of support that's intermediate in value).

¹⁷I've been concerned with transparency as a thesis that's supposed to help us account for self-knowledge. But slightly different ways of fleshing out the quote from Evans yield slightly different versions of transparency. Might some of these versions avoid at least some of the problems discussed in this section? Silins (2013), for instance, fleshes out Evans's remark by defending the following claim: When you judge that p , your judgement gives you immediate (and prima facie, propositional, and introspective) justification to believe that you believe p (p. 299). As it stands, the claim is about all-or-nothing belief. But a corresponding claim with respect to partial beliefs

4 Conclusion

Are partial beliefs transparent? I've clarified what it is to ask such a question. I've also considered several attempts to answer the question in the affirmative. But I've argued that such attempts face problems. So what?

As mentioned in section 1, it's possible that all-or-nothing beliefs are transparent whereas other attitudes—such as partial beliefs—are not. But the failure of transparency with respect to partial beliefs puts pressure on at least some philosophers, such as Byrne (2011), who would like transparency to give us a unified treatment of self-knowledge (p. 213). Also, we possess self-knowledge not only with respect to all-or-nothing beliefs, but with respect to partial beliefs. If one thinks that self-knowledge with respect to the former can be explained by transparency, there is at least some pressure to explain why self-knowledge with respect to partial beliefs is obtained via a different route, especially given that both all-or-nothing beliefs and partial beliefs are doxastic attitudes.

You might think that, for all I've said, there is still some notion of objective probability that will help us formulate a workable thesis of transparency for partial belief. But I've argued that the usual notions do not help. If you think there's some notion that does, the burden of proof is on you to show how so.¹⁸

may go as follows: When you judge that p is $x\%$ (objectively) likely to be true, your judgement gives you immediate justification to believe that you are $x\%$ confident that p . Since the claim, as it stands, is silent about cases in which you judge that the probability of p 's being true is ill-defined and silent about cases in which you judge that the probability of p 's being true is either 0 or 1 (without judging that it's 0 or judging that it's 1), it's not susceptible to some of the problems raised above. But such silence also means that we can't appeal to it to account for self-knowledge of partial beliefs in a relatively large swathe of cases.

¹⁸Thanks to the audience at the 1st Veritas Philosophy Conference at Yonsei University and the audience at the 2014 Australasian Association of Philosophy Conference for their useful feedback. Thanks also to Jens Christian Bjerring, Ben Blumson, Mark D'Cruz, Jay Garfield, John Holbo, Ole Koksvik, Peter Kung, Leon Leontyev, Neil Mehta, Michael Pelczar, Hsueh Qu, and Nico Silins for valuable discussion. Last but not least, I'm grateful to the anonymous referees for this journal for their exceptionally helpful comments.

References

- Alston, W. P. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. US: Cornell University Press.
- Armstrong, D. M. (1968). *A Materialist Theory of the Mind*. London and New York: Routledge.
- Armstrong, D. M. (1981). *The Nature of Mind and Other Essays*. Ithaca, New York: Cornell University Press.
- Audi, R. (1995). Memorial justification. *Philosophical Topics* 23, 31–45.
- Bernecker, S. (2008). *The Metaphysics of Memory*. New York: Springer.
- Byrne, A. (2005). Introspection. *Philosophical Topics* 33, 79–104.
- Byrne, A. (2011). Transparency, belief, intention. *Proceedings of the Aristotelian Society Supplementary Volume* 85, 133–64.
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Christensen, D. (2004). *Putting Logic in Its Place: Formal Constraints on Rational Belief*. New York: Oxford University Press.
- Davies, M. (2015). Knowledge (explicit, implicit and tacit): Philosophical aspects. In J. D. Wright (Ed.), *International Encyclopedia of the Social & Behavioral Sciences (Second Edition)*, pp. 74–90. Oxford: Elsevier Ltd.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Feldman, R. and E. Conee (2001). Internalism defended. *American Philosophical Quarterly* 38, 1–18.
- Frankish, K. (2009). Partial belief and flat-out belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*, pp. 75–93. Dordrecht: Springer.

- Gertler, B. (2011). Self-knowledge and the transparency of belief. In A. Hatzimoysis (Ed.), *Self-Knowledge*. Oxford University Press.
- Gertler, B. (2015). Self-knowledge. *The Stanford Encyclopedia of Philosophy* [online encyclopedia], Summer 2015 Edition, <<http://plato.stanford.edu/archives/sum2015/entries/self-knowledge/>>.
- Goldman, A. I. (1999). Internalism exposed. *The Journal of Philosophy* 96, 271–93.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. Cambridge, MA: The MIT Press.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability: Volume II*, pp. 263–93. Berkeley and Los Angeles, CA: University of California Press.
- Maher, P. (2006). The concept of inductive probability. *Erkenntnis* 65, 185–206.
- Moran, R. (2001). *Authority and Estrangement*. Princeton: Princeton University Press.
- Paul, S. K. (2012). How we know what we intend. *Philosophical Studies* 161, 327–46.
- Pollock, J. L. (2006). *Thinking about Acting: Logical Foundations for Rational Decision Making*. New York: Oxford University Press.
- Schaffer, J. (2007). Deterministic chance? *British Journal for the Philosophy of Science* 58, 113–40.
- Senor, T. (1993). Internalist foundationalism and the justification of memory belief. *Synthese* 94, 453–76.
- Silins, N. (2013). Introspection and inference. *Philosophical Studies* 163, 291–315.
- Stace, W. T. (1952). *Religion and the Modern Mind*. Philadelphia: Lippincott.
- Yalcin, S. (2007). Epistemic modals. *Mind* 116, 983–1026.