

Chilling out on epistemic rationality

A defense of imprecise credences (and other imprecise doxastic attitudes)

Miriam Schoenfield

Published online: 21 March 2012
© Springer Science+Business Media B.V. 2012

1 Introduction

A central question in epistemology is: which beliefs are *rational*? However, I intend to argue that the notion of rationality is, in fact, a conflation of two quite different notions: the notion of *what we ought to believe* and the notion of *what the evidence supports*. I also think that many of the problems that epistemologists face can be solved by distinguishing these two notions. I will say more about these notions later in the paper, but for now, the crucial point to note is that *what beliefs we ought to have* depends, in part, on our cognitive capacities, whereas *what beliefs the evidence supports* does not. In this paper, I will focus on how this distinction can be used to address some puzzling issues that arise with regard to the question of how precise our doxastic attitudes should be.

The problem that I will be addressing is an instance of a more general kind of problem that the distinction between *what we ought to believe* and *what the evidence supports* can solve. The kind of problem I have in mind is of the following form: (1) It can plausibly be argued that if a belief (or, more generally, a doxastic attitude) is rational, some very demanding condition C must be met. (2) We notice that a lot of our beliefs fails to meet condition C, and so we feel pressure to give up these beliefs. (3) We don't want to give up a lot of our beliefs!

Some of the demanding conditions that epistemologists defend are, I think, conditions on evidential support, but *not* conditions on what we ought to believe. Once we distinguish these two notions, we can explain what is plausible about very demanding conditions on rational belief without being worried by the threat that we may actually have to give up all of the beliefs that fail to meet this condition. The demanding condition that will be the focus of this paper is the condition that *all* of our doxastic attitudes ought to be extremely precise. According to this condition,

M. Schoenfield (✉)
Massachusetts Institute of Technology, Cambridge, USA
e-mail: miriams@mit.edu

our belief-like states should be entirely representable by a precise probability function which assigns to each proposition a precise credence (a real number between zero and one that represents our degree of confidence in that proposition). Here is the claim:

PRECISION: Our doxastic state ought to be representable by a precise probability function.

I am going to argue that PRECISION is false, and that we can respond to some arguments in its favor by appealing to the distinction between what attitudes we ought to have, and what attitudes the evidence supports. I will argue that, although the evidence always supports precise credences, it is not the case that we always ought to have them.

Here is the plan: In sect. 2, I will give a rough overview of some of the advantages and disadvantages of PRECISION and talk about the bearing of this claim on some important philosophical issues. In sects. 3 and 4, I will present two arguments: one against PRECISION and one in its favor. The argument *against* PRECISION will appeal to a phenomenon that I will call “insensitivity to mild evidential sweetening.” I will argue that it is permissible to be insensitive to mild evidential sweetening, but that PRECISION prohibits this, and so, PRECISION must be false. The argument *for* PRECISION will appeal to a principle along the lines of van Fraassen’s famous Reflection Principle, which says, roughly, that you should defer to the opinions of your future, more informed, self. The combination of these arguments, as you may imagine, will leave us in a rather uncomfortable position. In sect. 5, I will describe the distinction between what we ought to believe and what the evidence supports in more detail and also mention some applications of the distinction. The remainder of the paper will be devoted to showing how, armed with this distinction, we can resolve the issues that arise with regard to the question of how precise our doxastic attitudes ought to be. I will show that we can do justice to the kernel of truth in both the arguments for and against PRECISION by allowing that, although the evidence supports doxastic attitudes that are extremely precise, the attitudes that we ought to have, are mushy.

2 PRECISION: Some initial pros and cons

You may find the claim that we are rationally required to have precise credences in every proposition so preposterous as to not even be worth considering. Later, I will talk about some arguments for the claim, but to hold you over, I would like to begin by describing some of the initial motivations for thinking that our doxastic attitudes should be precise.

There are a number of possible answers to the question of how precise our doxastic attitudes should be. At one extreme, there is a very coarse-grained picture, according to which there are only three attitudes that one might take towards p : believing that p , suspending judgment on whether p , or disbelieving that p . The problem with this picture is that it seems that, sometimes, we get evidence that warrants degrees of confidence *in between* these attitudes. Maybe, given E , we

should be *pretty confident* that p but not full-out believe p . We could add in some more options, like being pretty confident that p , and being pretty confident that $\sim p$. But that still might not be enough. To see this, suppose that you have two coins that are about to be tossed: one which you know to be fair, and one which you know to have a 0.51 bias towards the heads side (that is, the chance of the coin landing heads is 51%). Seemingly, you should be more confident that the second coin will land heads than that the first one will. However, if we want to capture such small differences, it seems that three, five, or seventeen options will still not be enough. Perhaps, then, what we need is an *infinite* number of options.

The drive for increasingly fine grained attitudes towards propositions has culminated in the idea that, in fact, our entire doxastic state should be representable by a probability function. Such a function assigns to each proposition a real number between 0 and 1, where 1 represents full belief (certainty), 0 represents full disbelief, and the interval in between represents the continuous range of degrees of confidence that one might have. Bayesian epistemology has made a business of appealing to such probability functions in constructing a theory of rationality. According to Bayesians, rational agents have attitudes that can be represented using a probability function which obeys the axioms of probability, and they revise those attitudes by Bayesian conditionalization.

Nonetheless, you might think that, even if *sometimes* it is appropriate to have a very fine grained attitude towards a proposition, rationality does not require that *all* of our doxastic attitudes be precise credences. Sometimes, it seems, we can rationally have attitudes that are imprecise. Consider, for example, the proposition that, on average, 24 men in Bulgaria stand on their heads on Sundays (I will call this proposition “B”). Is there really some precise probability that I should assign to such a proposition? This seems unlikely.

Thinking about propositions like B makes PRECISION look extremely implausible. In fact, some people have found PRECISION so implausible, that they see the Bayesian’s commitment to PRECISION as sufficient grounds for rejecting Bayesianism as a theory of epistemic rationality.¹ It is worth noting, however, that a degrees-of-belief framework, and Bayesianism in particular, have a lot going for them. Bayesianism has managed to solve some disturbing epistemological puzzles and paradoxes, connects well to our theory of practical rationality, and has shed light on issues in philosophy of science.² For this reason, there have been attempts to construct a version of Bayesianism that can survive the rejection of PRECISION. (As Levi so poetically notes, those who “object to the numerical precision required by strict Bayesianism are accommodated in the ample bosom of Mother Bayes”).³ I am somewhat optimistic. However, as I will argue, rejecting PRECISION is trickier than one may have thought (especially for the Bayesian). In the next two sections, I will present an argument against PRECISION, and an argument for PRECISION. I will then argue that we can do justice to these competing considerations by allowing that, while the evidence supports precise credences, it is not the case that we ought to have them.

¹ See, for example, Sober (2002).

² See, for example, Earman (1992), Howson and Urbach (1993) and Strevens (2005).

³ Levi (1985, p. 392).

3 An argument against PRECISION

In this section, I will present an argument against PRECISION, and I will also describe a model that we can use to represent agents whose attitudes are not representable by a probability function. Before presenting the argument against PRECISION, I should note that it is certainly not the only one. Many people⁴ have argued against PRECISION, primarily on the basis of the consideration that, in many situations, it seems extremely implausible to suppose that there is a unique real number which is the credence one should have in some proposition (recall the proposition B). I am somewhat sympathetic to a version of these considerations and, at a later point in the paper, I will talk about them in more detail. But I think that more can be said about why PRECISION must be false, and the aim of this section is to bring out a peculiar consequence of PRECISION.

Here is the structure of the argument: I am going to argue that it is permissible for us to have a pattern of doxastic attitudes which I will call “insensitivity to mild evidential sweetening” (to be explained in a moment). As I will show, agents that are insensitive to mild evidential sweetening have attitudes that cannot be represented by a probability function. Since these attitudes are permissible, PRECISION must be false.

I will begin by defining insensitivity to mild evidential sweetening:

INSENSITIVITY TO MILD EVIDENTIAL SWEETENING

You are insensitive to mild evidential sweetening with regard to p if and only if:

- (a) You are no more confident in p than you are in $\sim p$ and you are no more confident in $\sim p$ than you are in p (in other words, you are agnostic about whether p)
- (b) There is some piece of evidence E which supports p more than it supports $\sim p$.
- (c) If you learn E , you are *still* no more confident in p than you are in $\sim p$.⁵⁶

If you are insensitive to mild evidential sweetening with regard to p , I will sometimes say that you “*have no clue* whether p .” The argument will go as follows:

Argument against PRECISION

- (1) It is permissible to be insensitive to mild evidential sweetening.
- (2) If we are insensitive to mild evidential sweetening, our attitudes cannot be represented by a probability function.
- (3) It is permissible to have attitudes that are not representable by a probability function. (1, 2)
- (4) PRECISION is false. (3)

⁴ See, for example, Levi (1974), Jeffrey (1983), Kaplan (1996), Joyce (2005, 2010).

⁵ I should also note that one could be insensitive to evidential sweetening with respect to two propositions p and q where q is not the negation of p , and much of what I say here will apply to such cases as well. However, for simplicity, I am going to restrict my discussion to cases of insensitivity to evidential sweetening to cases in which the two propositions in question are mutually exclusive.

⁶ For discussions of insensitivity to sweetening in the context of practical rationality and ethics see, for example, Chang (1997), Hare (2010), and Schoenfield (ms.)^a.

3.1 Defense of (1)

I will now defend the claim that it is permissible to be insensitive to mild evidential sweetening. My defense of this claim will focus on the following case:

DETECTIVE CONFUSO

You are a confused detective trying to figure out whether Smith or Jones committed the crime. You have an enormous body of evidence that you need to evaluate. Here is some of it: You know that 68 out of the 103 eyewitnesses claim that Smith did it but Jones' footprints were found at the crime scene. Smith has an alibi, and Jones doesn't. But Jones has a clear record while Smith has committed crimes in the past. The gun that killed the victim belonged to Smith. But the lie detector, which is accurate 71% percent of the time, suggests that Jones did it. After you have gotten all of this evidence, you have no idea who committed the crime. You are no more confident that Jones committed the crime than that Smith committed the crime, nor are you more confident that Smith committed the crime than that Jones committed the crime.

Let's suppose that you are fully confident that either Smith or Jones committed the crime, and that your agnosticism about who committed the crime is a reasonable attitude to take. Let's call the proposition that Smith committed the crime, "S," and the proposition that Jones committed the crime, "J." Note that since you are fully confident that either Smith or Jones committed the crime, J can be treated as the negation of S (if S is false, then J is true, and if S is true, then J is false).

Now imagine that, after considering all of this evidence, you learn a new fact: it turns out that there were actually 69 eyewitnesses (rather than 68) testifying that Smith did it. Does this make it the case that you should now be *more* confident in S than J? That, if you had to choose right now who to send to jail, it should be Smith? I think not. This extra piece of evidence does not seem like enough to tip the scales.

If you think that, now that you have learned that there were 69, rather than 68, eyewitnesses, you ought to be more confident in S than J, consider the following disposition that you (most likely) have. If you had known all along that 69 people testified against Smith, you *still* would have been no more confident in S than J and no more confident in J than S. This disposition of yours seems completely appropriate. But if you deem this disposition to be appropriate, you should not now let the additional eyewitness make you more confident in S than J. The fact that you learned about this eyewitness at a later time should not affect whether this information has sufficient evidential force to warrant more confidence in S than J. (I am appealing here to the commutativity of evidence principle: the order in which you learn the evidence does not make a difference to what your doxastic attitude ought to be on its basis).

In our case, you are insensitive to evidential sweetening with respect to S since you are no more confident in S than $\sim S$ (i.e. J), and no more confident in $\sim S$ (i.e. J) than S. The extra eyewitness supports S more than it supports $\sim S$, and yet despite learning about the extra eyewitness, you are no more confident in S than you are in $\sim S$ (i.e. J).

3.2 Defense of (2)

I will now argue that, if you are insensitive to evidential sweetening with regard to p , your attitude towards p cannot be represented by a precise credence. The basic idea is this: precise credences are real numbers, and real numbers have the following property: for real numbers r_1 and r_2 , if r_1 is no greater than r_2 , and r_2 is no greater than r_1 , then $r_1 = r_2$. Now suppose that you have no clue whether p . Then, since you are no more confident in p than in $\sim p$ and no more confident in $\sim p$ than in p , if your attitudes can be represented by real numbers, your credences in p and $\sim p$ must equal the same real number. But if they are equal, and then we raise your confidence ever so slightly in p , the scales will be tipped: your new credence in p will now be greater than your credence in $\sim p$. However, if you are insensitive to evidential sweetening, getting a bit of evidence for p does *not* make you more confident in p than in $\sim p$. So, if you are insensitive to sweetening, your doxastic attitudes cannot be represented by real numbers.

It is important to note that, unlike the case of Detective Confuso, there are many cases in which you are no more confident in p than you are in $\sim p$ and no more confident in $\sim p$ than you are in p , and yet you *are* sensitive to evidential sweetening. For example, suppose you have a coin which you take to be fair. Let H be the proposition that the coin will land heads, and T be the proposition that the coin will land tails. Since you believe that the coin is fair, you are no more confident in H than T and no more confident in T than H . If you get any evidence that the coin is even *slightly* weighted towards the heads side, you will become *slightly* more confident in H than you are in T . (Note that, unlike in the case of S and J , you would be disposed to be more confident in H than T even if you had known that the coin was biased towards heads all along). What we can learn from these examples is that there are different ways in which you can be no more confident in p than $\sim p$ and no more confident in $\sim p$ than p . You might *have no clue whether p* (in which case you will be insensitive to sweetening) or you might have *equal credence in p and $\sim p$* (in which case you will be sensitive to sweetening). In the former case, your attitudes, though reasonable, cannot be represented by real numbers. So PRECISION is false.

3.3 Representing insensitivity to sweetening

Although we cannot represent an agent's doxastic attitudes by real numbers if that agent is insensitive to sweetening, we *can* represent such attitudes with a set of probability functions, called a *representor*.⁷ Here is how this will work: when you are more confident in p than $\sim p$, all of the probability functions in your representor will assign a greater credence to p than to $\sim p$. If you are equally confident in p and $\sim p$, all of the functions in your representor will assign credence 0.5 to p and 0.5 to $\sim p$. If you have no clue whether p or $\sim p$, the functions in your representor will differ in the relative credences assigned to p and $\sim p$ (for example, some may assign higher credence to p than to $\sim p$, while others assign a higher credence to $\sim p$ than to p).

⁷ For early discussions of this model see Jeffrey (1983) and Levi (1985).

The reason this structure allows for insensitivity to sweetening is as follows: You might start out no more confident in p than $\sim p$ and no more confident in $\sim p$ than p . This could happen if some of the functions in your representor assign a higher credence to p than to $\sim p$ and some assign a higher credence to $\sim p$ than to p . You then get some small bit of evidence for p . As a result, all of the functions in your representor will now assign a slightly higher credence to p than they did before. However, this does not imply that all of the functions will now assign a higher credence to p than to $\sim p$. Thus, you can be more confident in p than you were before (and hence, responsive to your new evidence) without becoming more confident in p than $\sim p$. We can use this model, then, to represent the phenomenon of insensitivity to sweetening. This will be useful later on.

4 An argument for PRECISION

We have seen that there is good reason to reject PRECISION: namely, that PRECISION is inconsistent with the claim that we should sometimes be insensitive to mild evidential sweetening. There is, however, a compelling argument in its favor.⁸ The gist of the argument is that if you violate PRECISION, because you are insensitive to mild evidential sweetening, it seems that you will violate the following plausible principle:

REFLECTION: If you know that, in the future, you will rationally have doxastic attitude A towards p , without any loss of information, you ought to now have doxastic attitude A towards p .⁹

This principle is in the same spirit as the REFLECTION principle that was first introduced and defended by van Fraassen (1984) and, for reasons of space, I will not venture into an extended defense of REFLECTION here. But the intuitive idea behind REFLECTION is this: From an epistemic standpoint, more evidence is better. Since you know that your later self will be better informed, and that the judgment made on the basis of that additional information is the rational judgment, you should view your later self as an expert (at very least, an expert relative to your current self). Thus, if you know what your later, more informed, and rational attitude towards p is, you should adopt that attitude now.

To see why subjects who are insensitive to evidential sweetening seem to face a violation of REFLECTION, I will describe a different version of the detective case, which I will call DETECTIVE CONFUSO–OPAQUE. I will argue that, in this version, you

⁸ In fact, there are a number of compelling arguments for PRECISION (see Elga 2010; White 2010) and I won't be able to address all of them here. The response I give to the argument that I will be presenting is also responsive to White's argument for PRECISION, and I think it may have some bearing on Elga's argument as well, but I will leave that for another time.

⁹ A more precise version of the principle would say that if you are fully confident that your future doxastic attitude will be A , you should now adopt A , and if you are less than fully confident that your future attitude will be A , your attitude should be an average of the possible attitudes you might have, weighted by the probability of you having those attitudes. For our purposes, however, this rough version is good enough.

should be sensitive to evidential sweetening with respect to a proposition p . However, although you should be more confident in p than $\sim p$ after getting some small piece of evidence for p , you will know that, later, upon gaining more information, you will have no clue whether p . Thus, what your attitude should be now differs from what you know your future, more informed attitude will be.

Here is the case:

DETECTIVE CONFUSO—OPAQUE

As above, you have examined lots of evidence with regard to a crime committed by Smith or Jones and you have no clue as to whether Smith or Jones committed the crime. Smith and Jones were each placed in one of two jail cells—Cell #1 and #2. Who was placed in which cell was determined by the flip of a fair coin (if the coin landed heads, Smith is in Cell #1 and Jones is in Cell #2, and if the coin landed tails, Jones is in Cell #1 and Smith is in Cell #2). You don't know how the coin landed and you won't find out who was placed in which cell until tomorrow. In the meantime, an eyewitness comes into the police department, looks at the prisoners in each cell, and testifies that the person in Cell #1 committed the crime.

Call the proposition that the person in Cell #1 committed the crime “Cell 1” and the proposition that the person in Cell #2 committed the crime “Cell 2.” I will outline the argument for *PRECISION* below and then defend each of the premises.

Argument for PRECISION

1. In the opaque case, you should be more confident in Cell 1 than in Cell 2.
2. If you violate *PRECISION*, then, in the opaque case, you know that tomorrow, once you find out who is in which cell, you will rationally have no clue whether Cell 1 or Cell 2.
3. If you violate *PRECISION*, you violate *REFLECTION*. (1, 2)

Since you should not violate *REFLECTION*,

4. *PRECISION* is true. (3)

4.1 Defense of (1)

The first step is to defend the claim that, in *DETECTIVE CONFUSO—OPAQUE*, you should be more confident in Cell 1 than Cell 2 when you learn about the additional eyewitness. I will present three arguments for this claim:

4.1.1 *Argument #1: The reasons for belief argument*

What reasons do you have to believe Cell 2? Since you don't know who is in which cell, your only reason to believe that the person in Cell #2 committed the crime is that you know that the person in Cell #2 is either Smith or Jones, and that there is a body of evidence that you have examined which includes incriminating evidence against both of them. What reasons do you have to believe Cell 1? Here too, you

only know that the person in Cell #1 is either Smith or Jones and that there is a body of evidence you have examined which includes incriminating evidence against both of them. Thus, the reasons to believe Cell 2 are also reasons to believe Cell 1. However, you have an *additional* reason to believe Cell 1 which is not a reason to believe Cell 2, namely that the eyewitness claims to have seen the person in Cell #1 commit the crime. Since every reason you have to believe Cell 2 is also a reason to believe Cell 1, but you have an additional reason to believe Cell 1 that is not a reason to believe Cell 2, it follows that you have more reason to believe Cell 1 than Cell 2.¹⁰

4.1.2 Argument #2: The representor argument

The second argument for the claim that you ought to be sensitive to evidential sweetening in the opaque case appeals to the representor model I discussed previously. If this kind of approach is correct, it will follow that you ought to be more confident in Cell 1 than Cell 2. To see why, recall that if all of your attitudes *can* be represented by precise credences, then for any two propositions such that you are no more confident in the one than the other and no more confident in the other than the one, you will be sensitive to mild evidential sweetening. Now, once Smith and Jones are placed in the jail cells and you don't know who is in which cell, all the probability functions in your representor will assign 0.5 credence to Cell 1 and 0.5 credence to Cell 2.¹¹ But, since each of these functions is a precise probability function, each function will, individually, be sensitive to mild evidential sweetening. So, once you learn about the eyewitness, all of the functions in your representor will assign higher credence to Cell 1 than Cell 2 (and thus *you* are more confident in Cell 1 than Cell 2).

4.1.3 Argument #3: The argument from $Pr(\text{heads}) = 0.5$

Let's consider a slight variant of the case. As before, we imagine that who was placed in which cell was determined by the flip of a fair coin. But this time, we assume that the warden knows which of the two suspects is guilty and that the way the coin flip worked is as follows: If the coin lands heads, the warden puts the guilty person in Cell #1 and the innocent person in Cell #2. If the coin lands tails, the warden puts the innocent person in Cell #1, and the guilty person in Cell #2.

¹⁰ This argument was inspired by a similar argument in decision theory described in Hare (2010).

¹¹ To see why, recall that who was put in which cell was determined by the flip of a fair coin. If the coin landed heads (H), Smith is in Cell #1 and if the coin landed tails (T), Jones is in Cell #1.

- (1) $\text{Cell 1} \leftrightarrow [(S \& H) \text{ or } (J \& T)]$ (in other words, the person in Cell 1 is guilty if and only if Smith is guilty and the coin landed heads or Jones is guilty and the coin landed tails).
- (2) $\text{Pr}(\text{Cell 1}) = \text{Pr}[(S \& H) \text{ or } (J \& T)]$
- (3) $\text{Pr}(\text{Cell 1}) = .5 \text{Pr}(S) + 0.5 \text{Pr}(J)$
- (4) $\text{Pr}(J) = 1 - \text{Pr}(S)$ (since either Smith or Jones is guilty)
- (5) $\text{Pr}(\text{Cell 1}) = 0.5 \text{Pr}(S) + 0.5 [1 - \text{Pr}(S)] = 0.5$
- (6) $\text{Pr}(\text{Cell 2}) = 1 - \text{Pr}(\text{Cell 1}) = 0.5$.

	Cell #1	Cell #2
Coin lands heads →	Guilty	Innocent
Coin lands tails →	Innocent	Guilty

Recall that Cell 1 is the proposition that the person in Cell #1 is guilty, and let H be the proposition that the coin lands heads. Once you know how the coin flip works, you know that Cell 1 is true if and only if H is true. Since it is clear that you should have a 0.5 credence in H , and since your attitude towards H and Cell 1 must be the same, you should have a 0.5 credence in Cell 1 as well.¹² But if you have a precise credence in Cell 1, you will be sensitive to mild evidential sweetening, and so, when the eyewitness comes in, you will be more confident in Cell 1 than Cell 2. Thus, at very least, in this version of the opaque case, you should be more confident in Cell 1 than Cell 2.¹³ This is all we will need to get the problem with REFLECTION going.

4.2 Defense of (2)

I have described three arguments which support the claim that, in the opaque detective case, you ought to be sensitive to evidential sweetening; that is, after the eyewitness comes in, you ought to be slightly more confident that the person in Cell #1 committed the crime than that the person in Cell #2 committed the crime. I will now argue that you know that, in the future, you will have no clue whether Cell 1 or Cell 2. To see this, consider what happens, in the future, when you learn which suspect is in which cell.

First, suppose you learn that Smith is in Cell #1. If you learn that Smith is in Cell #1, your doxastic attitude towards Cell 1 will be the same as your doxastic attitude towards S . What will your doxastic attitude towards S be? Since you are insensitive to evidential sweetening with regard to S , even though you will know that an additional eyewitness testified against Smith, you will still be no more confident in S than J and no more confident in J than S . And since you know that Smith is in Cell #1 and Jones is in Cell #2, you will also be no more confident in Cell 1 than Cell 2 and no more confident in Cell 2 than Cell 1.

Now, suppose you learn that Jones is in Cell #1. Since you are insensitive to evidential sweetening with regard to J , even though you will know that an additional eyewitness testified against Jones, you will be no more confident in J than S , and no more confident in S than J . Thus, you will also be no more confident in Cell 1 than Cell 2 and no more confident in Cell 2 than Cell 1. So you know that, *no matter how things turn out*, you will later have no clue whether Cell 1 or Cell 2. But then it seems that, according to REFLECTION, you ought to not have a clue whether Cell 1 or Cell 2 *even before you find out who is in which cell*, contrary to the judgment that, in the opaque case, you ought to be more confident in Cell 1 than Cell 2.¹⁴ In the next

¹² There is an argument in White (2010 pp. 175–181), which, in this case, could be applied to defend the claim that you should match your credence in Cell 1 to your credence in H , rather than the other way around.

¹³ This way of motivating the claim that, in the opaque case, you should be sensitive to sweetening, was inspired by a case discussed in White (2010).

¹⁴ Some proponents of imprecise credences might think that the correct version of the reflection principle will only tell you to defer to your future doxastic states if you know what your *entire representor* will be

section, I am going to describe a strategy which will allow us to have imprecise attitudes without facing a reflection violation.

5 Distinguishing between what we ought to believe and what the evidence supports

In order to solve our problem, I am going to need to appeal to a distinction between what doxastic attitudes the evidence supports and what doxastic attitudes we ought to have. One way to get a feel for the distinction is to think of it as a distinction between what attitudes agents with perfect cognitive capacities would have, and what attitudes agents like us, with various limitations, ought to have.¹⁵ I defend this distinction in greater detail elsewhere¹⁶ and, although I cannot provide the full argument for the distinction here, I think it is rather intuitive, so I will just say a few things which I hope will convince you of its importance.

To get started, I am going to make some assumptions about the evidential support relation. The first assumption is that the evidential support relation has the following feature: if the evidence entails *p*, it supports a high degree of confidence in *p*.¹⁷ The second assumption is that the evidential support relation is a relation that holds between bodies of evidence and doxastic attitudes. Each body of evidence supports a doxastic attitude and which attitudes a body of evidence supports is determined completely by features of the evidence. In particular, what the evidence supports does not depend on the particular agent who is evaluating the evidence.¹⁸

Footnote 14 continued

at the later time. This condition is not satisfied in this case. However, I think it would be a mistake to restrict reflection principles in this way. We don't want the principles that tell us how to defer to experts (whether they are other people, or just future time slices of ourselves) to be applicable only in cases when we know what the expert's entire representor is, since we rarely have such information.

¹⁵ More precisely, it is a distinction between what attitudes we should have, and what attitudes agents with perfect cognitive capacities, and who are *unreflective*, would have—where by this I mean, that these agents don't worry about the possibility of their own error. The addition of the “unreflectivity” requirement is important for reasons discussed by Christensen (2008), and is necessary for agent neutrality, which will be discussed shortly. (This kind of perfect rationality is related to the notion Hartry Field (2000) describes as “ideal credibility”). For convenience, in what follows, I will use the term “agents with perfect cognitive capacities” to refer to unreflective agents with perfect cognitive capacities. Since I do not think that any agents should be unreflective I hesitate to use this terminology. I use it anyway, since I think it conveys something of the idea I'm trying to develop. A less agent-centered (and perhaps entirely uninformative) way of thinking about the degree of confidence in a proposition that the evidence supports is as its evidential probability.

¹⁶ Schoenfield (ms.)^b. Also, see Aarnio (2010) and Sepielli (ms.) for discussions of distinctions along these general lines.

¹⁷ I do not mean to suggest that all evidence is propositional, but only that, for those propositions that are part of our evidence if they entail *p*, then our evidence supports a high degree of confidence in *p*.

¹⁸ Two notes here: First, the agent neutrality condition applies to *de dicto* propositions only. Second, if you are a permissivist, and think that what *S*'s evidence supports depends on *S*'s priors, or standards of reasoning, we can let the evidential support relation be a three place relation between the evidence, the agent's priors, and doxastic attitudes. It will still be true that what the evidence supports does not depend on which *particular* agent is evaluating it (though what the evidence supports will depend on the agent's priors, or standards of reasoning).

With these assumptions in hand, I am now going to provide a sketch of the argument I use to defend the claim that the attitudes supported by the evidence may be different from the attitudes we ought to have. First, since the evidence always supports being confident in those propositions that it entails, and all mathematical truths are entailed by any body of evidence, if we ought to have the attitudes supported by the evidence, then we ought to be confident in all the mathematical truths. But it is not true that we should be confident in all of the mathematical truths—not only because we could not—but because the claim that we should be confident in all of the mathematical truths conflicts with a variety of plausible principles about what attitudes we should have towards mathematical propositions. For example, it is plausible that there are some complex mathematical propositions that we should suspend judgment on (like the proposition that the millionth digit of π is even). But this conflicts with the claim that we should be confident in all true mathematical propositions. It is also plausible that we should defer to the experts about some mathematical propositions (for example, Fermat's Last Theorem). Even if Fermat's Last Theorem turns out to be false, it would still be true that we should now be confident in its truth, because of the experts' testimony. However, if we should be confident in all mathematical truths, then, if Fermat's Last Theorem turns out to be false, we should not be confident in it now, despite the fact that the experts say we should. Thus, it cannot be right that we should be confident in all of the mathematical truths, and so it cannot be right that we should always be confident in what our evidence entails.

Another motivation for the claim that we should not always have the attitudes best supported by the evidence comes from considerations of higher order evidence. Suppose that you are evaluating some proposition p , and come to believe p , when you realize that you might be under the influence of a reason distorting drug. In such a case, you should reduce confidence in p . However, your friend, who shares all of your evidence,¹⁹ but has no reason to think that *she* is under the influence of a drug, should not reduce her confidence in p .²⁰ Since I am assuming that what the evidence supports does not depend on which particular agent is evaluating it, the evidence cannot support your friend believing p , and you reducing confidence in p . Nonetheless, your friend should be confident in p and you should not. So here too we have a case in which what the evidence supports comes apart from what you ought to believe.

These are the sorts of considerations that support the claim that we should not always have the attitudes best supported by the evidence. Now, you might think that it would be better to maintain that we ought to always have the attitudes supported by the evidence and change our views about what features the evidential support relation has. Maybe what these considerations show is that the evidence does *not* always support what it entails, and perhaps the evidential support relation *is* an agent relative one. May be. But, as I will illustrate, the notion of evidential support that I described is going to play an important role in our theory, and tinkering with it

¹⁹ If you are a permissivist, we will add to this the qualification that your friend has the same standards of reasoning, or prior probability function as you do.

²⁰ This phenomenon was first discussed (as far as I know) in Christensen (2010).

is a bad idea. For, even if it is not always the case that the doxastic attitudes best supported by the evidence are the ones we ought to have, what the evidence supports may turn out to be highly relevant to the question of which attitudes we ought to have.

The distinction between what the evidence supports and what we ought to believe is crucial to the solution of our puzzle. For convenience, I am going to say that some doxastic attitude is *reasonable* for you, given some body of evidence, if that is the attitude you ought to have given that body of evidence. Here is not the place to go into a full exploration of the distinction between what is reasonable and what is supported by the evidence. In fact, most of what follows does not rest much on the particular way we understand this distinction. But, since you may be wondering, I would like to say a bit about more about what I think being reasonable is all about.

Here is a rough sketch of what I take reasonability to be: I think reasonability is a property of doxastic attitudes that plays an important role in the *deliberative* perspective. This is because I think that the principles that determine what is reasonable are principles of deliberation that the agent has good reason to expect have the following feature: deliberating using these principles will help her with her epistemic aims (believing truths and avoiding falsehoods, or, if you prefer, maximizing her expected accuracy²¹).²² Some modifications and refinements of this condition need to be made to make it more precise and avoid some straightforward objections. But, to get a general sense of how it works, consider one of our examples. Why is it reasonable to defer to the mathematical experts about whether some mathematical proposition *M* is true, even though, as a matter of fact, $\sim M$ is entailed by your evidence? Because, I claim, the principle: “defer to the mathematical experts” is a principle we have good reason to expect is such that deliberating with it will help us achieve our epistemic aims.

Evidential support, in contrast, is a property that, I think, is most usefully utilized in the *evaluative* perspective. The reason it is important to evaluate the way other agents reason is that a large part of our epistemic life involves deferring to other agents, and, as I will argue later, whether or not we defer to an agent has more to do with whether the agents’ attitudes are supported by the evidence than whether or not they are reasonable. The fact that we use evidential support when evaluating other agents is not inconsistent with the claim that what we ought to aim for is reasonability—for due to our own limitations, *trying* to be reasonable will get us further than *trying* to have the attitudes that the evidence supports.

Here is another way of putting the point: When deciding whether to defer to someone, it can make sense to see how they measure up against an ideal that may be incredibly difficult to satisfy. But, when deciding what to believe, it may be a bad

²¹ There are different ways of measuring accuracy, but the general idea is that an accurate agent will have high credences in truths and low credences in falsehoods.

²² Note that deliberating with a principle does not require successfully following it. If the principles of reasonability were the ones such that successfully *following* them would help us achieve our epistemic aims—the only principle we would need would be one which told us to be fully confident in all and only the truths. Since you can be reasonable without being fully confident in all and only the truths, it is important that the test for a principle of reasonability be concerned with the result of *trying* to follow the principle rather than the result of following it.

idea to do so by trying to satisfy principles that are incredibly difficult. This is because there may be alternative principles which would be more helpful. This is why there are not many constraints on how demanding the principles of evidential support can be, but there are significant constraints on how demanding the principles of reasonability can be.

In what follows, I will spell out how this distinction can solve our puzzle, by showing that if we accept the following combination of views we will be problem free:

- (a) The evidence supports precise credences.
- (b) It's not the case that we ought to have precise credences.

The reason that this combination of views will solve the problem is as follows: (1) Even if we grant that *the evidence supports* precise credences, it can still be true that we *ought* to be insensitive to mild evidential sweetening. (2) Granting that we ought to have imprecise attitudes, even though the evidence does not support imprecise attitudes, will not result in a reflection violation (this will be argued for in sect. 7). This is how the competing pressures concerning precise credences will get resolved. I will spell out the solution in more detail in sect. 7, but in the next section, I will describe some reasons why some people might not find this solution entirely satisfying.

6 An opportunity to get off the boat (and some encouragement to stay on)

In order to solve the problem, we are going to need to grant that, although *we* should not have precise credences, the evidence does indeed support precise credences. That the evidence supports precise credences is a crucial part of the solution, and, for this reason, my solution may not satisfy all defenders of imprecise attitudes. Some philosophers have thought that, not only are we not *required* to have precise credences, but that the evidence just *cannot* always support a precise credence. (If this is not a concern of yours, feel free to skip to sect. 7, in which I spell out the solution). Joyce (2005, p. 171) writes as follows:

...the proper response to symmetrically ambiguous or incomplete evidence is not to assign probabilities symmetrically, but to refrain from assigning precise probabilities at all...It is not just that sharp degrees of belief are psychologically unrealistic (though they are). Imprecise credences have a clear epistemological motivation: they are the proper response to unspecific evidence.

And Levi (1985, p. 396) writes:

...it should be emphasized that those who insist on the reasonableness of indeterminacy in probability judgment under the permissibility interpretation mean to claim that even superhumans ought not always to have credal states that are strictly Bayesian.

People like Joyce and Levi might claim that, in conceding that the evidence supports precise credences, I am not fully respecting the considerations that motivated the claim that we should have imprecise attitudes to begin with. (Though note that, in claiming that the evidence supports precise credences, I am still respecting the motivation for imprecise attitudes that I provided, which was that sometimes we ought to be insensitive to evidential sweetening).

While I am fully convinced that precise credences are not the sorts of things that we should always have, it is not obvious to me that the evidence cannot support precise credences. There are a number of reasons for this: First, I am not fully convinced by the motivations provided by Joyce and others; second, I am not sure that the problem which they are concerned with is really a problem that can be solved by moving from precise credences to imprecise credences; and finally, I think that some of these motivations rely on an assumption which I deny. Let me elaborate briefly on each of these of these considerations.

First, Joyce and others talk about “incomplete” evidence, which makes it sound like, sometimes, there simply is not *enough* evidence to support a precise credence. Propositions like, “On average, 24 men in Bulgaria stand on their heads on Sundays” are meant to illustrate this fact. But it does not seem like the problem here is a *lack* of evidence. After all, I know all sorts of things about the kind of circumstances which lead people to stand on their heads and I have some ideas about the population of Bulgaria, the prevalence of yoga classes in Europe, etcetera. In fact, I have *lots* of evidence relevant to the question of how many men in Bulgaria are now standing on their heads! The problem, then, is not that I don’t have enough evidence, but that the evidence is *complicated*. And it does not seem that the evidence being complicated is sufficient grounds for thinking that it does not support a precise credence.

Even if the standard cases that are appealed to are not ones in which the problem is insufficient evidence, we might be able to imagine such cases (though they will be quite unrealistic). For example, suppose you come into existence and all you see in front of you is a picture of an elephant. You then think to yourself, “I wonder if there exist at least seventeen creatures of this kind.” This case does indeed seem like a case in which the problem is an insufficiency of evidence. And, in cases like this, many of the proponents of imprecise credences would say that your attitude should be represented by a credence spread over the entire $[0, 1]$ interval. This judgment, however, seems to me to be unacceptable. Aside from the fact that there are problems with updating when one starts out in such a dilated state,²³ it seems to me that a fully rational agent’s attitudes will have significantly more structure than this suggestion supposes. For example, the fully rational agent will, I think, be more confident in the proposition that there are at least seventeen four legged creatures than in the proposition that there are at least seventeen elephants. So it will not do, in cases of insufficient evidence, to take the same attitude towards every proposition we do not have much evidence for. This means that, even if we allow that the fully rational agent would have degrees of confidence that could only be represented by

²³ At least according to Bayesian conditionalization, once a function assigns 1 or 0 to a proposition, it will always assign 1 or 0, no matter how much new evidence one conditionalizes on.

intervals, these intervals would be varied.²⁴ Perhaps the fully rational agent would have credence (0.2–0.3) towards the proposition about elephants, but (0.4–0.5) towards the proposition about four legged creatures. This brings me to the second worry I have with some of the motivations that have been given for imprecise credences. Whatever strangeness there seems to be in judging that fully rational agents would have precise credences in cases of insufficient evidence does not seem to me to disappear by allowing that these agents would have attitudes representable by intervals. I think that being fully rational requires that, even in cases in which there is not much evidence, one's degrees of confidence be structured in various ways, and it does not seem to me that there is a significant difference in the intuitive plausibility of a structure represented by numbers, and one represented by intervals.

Finally, I think that some of the motivations for the claim that the evidence doesn't support precise credences get their bite through reliance on a *uniqueness* assumption; that is, they assume that, if the evidence supports having a doxastic state that is representable by a precise credence function, there is only one precise doxastic state that is warranted given a particular body of evidence. And indeed, the thought that, given some body of evidence, there is a precise credence that the evidence supports towards any given proposition can seem "spooky." However, nothing that I have said here precludes the possibility that there are *multiple* precise credences one could have that are consistent with being perfectly rational. That is, although I am claiming that agents with doxastic states supported by the evidence will have precise credences, I am not claiming that, given any body of evidence, there is a unique precise credence that the evidence supports (in fact, I think the *uniqueness* assumption is false, but I will not get into the details of that here).²⁵

7 The solution

Here is where we are so far: I have presented arguments both for and against PRECISION and I have suggested that we can resolve these competing pressures by adopting the following combination of views: (a) The evidence supports precise credences, and (b) it's not the case that we ought to have precise credences. Granting that the evidence supports precise credences still allows us to be insensitive to mild evidential sweetening, and, as I will show here, granting that we are permitted to have imprecise attitudes does not result in a reflection violation.

Before explaining why the combination of views mentioned above is not subject to a reflection worry, let me first remind you of the principle in question:

REFLECTION: If you know that, in the future, you will *rationaly* have doxastic attitude A towards p, without any loss of information, you ought to now have doxastic attitude A towards p.

²⁴ Some people (Joyce 2010 for example) might object by saying that the structure does not need to show up in the *intervals* that represent the agent's attitudes towards individual propositions, so long as the structure is found in the representor *as a whole*. However, I think it is important that we be able to represent an agent's attitude towards a single proposition without building in information about the agent's entire representor. (One reason for this is described in footnote 13).

²⁵ I argue against the *uniqueness* assumption in Schoenfield (ms)^c.

REFLECTION says that you should defer to your future attitudes if those attitudes are rational. Now that we have distinguished the attitudes that are supported by the evidence from the attitudes that are reasonable, we can ask, which of these does “rational” in REFLECTION refer to. In other words, which of the following principles are true:

EVIDENTIAL REFLECTION: If you know that, in the future, you will have doxastic attitude *A* towards *p*, and that *A will be the attitude best supported by the evidence*, then (provided you lose no information), you ought to now have doxastic attitude *A* towards *p*.

REASONABLE REFLECTION: If you know that, in the future, you will have doxastic attitude *A* towards *p*, and that *A will be an attitude that is reasonable given the evidence*, then (provided you lose no information), you ought to now have doxastic attitude *A* towards *p*.

I am going to argue that EVIDENTIAL REFLECTION is true, and that REASONABLE REFLECTION is false. But before doing so, let me explain how this will help with our problem. In the detective case, you know that, in the future, you will not have a clue whether Cell 1 or Cell 2, but you also know that this attitude cannot be represented by a precise probability. Since your later attitude is imprecise, that attitude is not one best supported by the evidence, and you do not violate EVIDENTIAL REFLECTION by failing to defer to it. It does not matter that you violate REASONABLE REFLECTION, since this principle is false anyway.

While the suggestion above may seem like a sneaky way of getting out of the problem, it is actually quite commonsensical. Reflection principles are based on the idea that you should defer to your later attitudes when you know that your later epistemic state is no worse than your current one. It is usually thought that, to be epistemically worse off, you must either lose information or become less capable of evaluating evidence. In this case, you do not lose information, nor do you lose any capacities. What does change, however, is the complexity of the evidence that you have to take into account. Before you know who is in which cell, the evidence that you have to evaluate is very simple. However, once you know who is in which cell, to determine what attitude the evidence supports having towards Cell 1, you need to be able to take account of a complex body of evidence in an incredibly precise way. Because of the complexity of the evidence you will have to evaluate, when you learn who is in which cell, your judgment will not be one that is supported by the evidence, and so you do not need to defer to it.

7.1 Why we should reject REASONABLE REFLECTION

The solution above works because REFLECTION only applies when one takes one’s future doxastic attitude to be one that is *supported by the evidence*. In this subsection, I will describe in greater detail why this is the case, and why REASONABLE REFLECTION, a principle that tells you to defer to future *reasonable* attitudes, is false. Recall REASONABLE REFLECTION:

REASONABLE REFLECTION: If you know that, in the future, you will have doxastic attitude A towards p, and that A *will be an attitude that is reasonable given the evidence*, then (provided you lose no information), you ought to now have doxastic attitude A towards p.

If this principle were correct, the puzzle would remain unsolved. This is because, while you do not take your future not-having-a-clue attitude to be supported by the evidence, you do take it to be reasonable. It would follow, then, from REASONABLE REFLECTION, that you must defer to this later attitude, and thus, that you should adopt the not-having-a-clue attitude towards Cell 1 even before you learn who is in which cell, contrary to our assumption. Fortunately, there are good reasons to reject REASONABLE REFLECTION.

7.1.1 Argument #1 for rejecting REASONABLE REFLECTION

The first reason to reject REASONABLE REFLECTION is that it is unmotivated. Recall that reflection principles are motivated by the idea that you think of your future self as an expert—that is, you should think of your future self as being in a better position to form true beliefs and avoid false ones than your current self. Now, if your future self's attitude is supported by the evidence, and she has more information than you do, it makes sense to suppose that she is in a better position than you are to form true beliefs and avoid false ones. If your future self's attitude is merely reasonable, however, then it's not always going to be true that she is in a better position than you are to form true beliefs and avoid false ones. This is because how likely you are to form true beliefs depends *both* on how much evidence you have and how well that evidence is evaluated. In the case in which your future self is merely reasonable, you know that your future self will have more information than you do, but your current self might be more likely to have the attitudes supported by the evidence than your future self (even though your future self is reasonable). So, the mere fact that your future self will be reasonable and have more information than you do, does not guarantee that she is better posed for forming true beliefs and avoiding false ones.

7.1.2 Argument #2 for Rejecting REASONABLE REFLECTION

The second reason to reject REASONABLE REFLECTION appeals to the following principle, which bridges what the evidence supports and what is reasonable.

DEFERENCE: If you know that the only doxastic attitude towards p which is supported by your evidence is A, reasonability requires that you have doxastic attitude A towards p.

The motivation for DEFERENCE is simple: having attitudes that are best supported by the evidence is, in a sense, better than being merely reasonable. (It is better in the sense that agents who always have the attitudes supported by the evidence will probably have more true beliefs and fewer false ones than merely reasonable agents, or, in a degrees of belief framework, it is likely that agents with the attitudes

supported by the evidence will be more accurate than those who do not.) The reason it is not the case that *we* should always have the attitudes best supported by the evidence has to do with facts about our cognitive limitations (or facts about the possibility of such limitations). However, when we are lucky enough to know of an attitude that it is the unique attitude supported by the evidence, surely that is the attitude we should have!

The problem with REASONABLE REFLECTION is as follows: If we accept REASONABLE REFLECTION, there is the potential to run into situations in which we know what attitude the current body of evidence supports, but we also know that, in the future, we will have some merely reasonable attitude which differs from the one that the current body of evidence supports. REASONABLE REFLECTION would then yield the result that we should have an attitude different from the one we know to be supported by the evidence, and so REASONABLE REFLECTION would sometimes have us violate DEFERENCE.

We can actually use the very case we have been considering to illustrate how REASONABLE REFLECTION leads to a violation of DEFERENCE. I will do so by showing that, in DETECTIVE CONFUSO—OPAQUE, you know that being more confident in Cell 1 than Cell 2 is the attitude supported by the evidence, and yet, REASONABLE REFLECTION would say that you should have no clue whether Cell 1 or 2. Thus, REASONABLE REFLECTION would tell you to have an attitude which differs from the one you know to be supported by the evidence.

Why do you know that being more confident in Cell 1, in the opaque case, is the attitude that is best supported by the evidence? Consider an agent who always has the attitudes supported by the evidence. Call this agent, “Alice.” How would Alice respond to your conundrum in the opaque detective case, were she to have the same evidence that you do? Note that Alice would not know what *her* future attitude would be because her future attitude towards Cell 1 would depend on who turns out to be in Cell 1. If, for example, she were slightly more confident in S than in J, and it turned out that Smith was in Cell #1, she would become more confident in Cell 1 than Cell 2. This might not be the case, however, if Smith turned out to be Cell #2. Even though Alice does not know what her future attitude will be, we may ask, how would Alice respond to learning that *you*, in the future, will have no clue as to whether Cell 1. Would this affect *her* confidence in Cell 1?

The answer is no. In general, Alice will not always defer to your more informed judgments since your judgments are not always supported by the evidence. Furthermore, since you will have no clue whether Cell 1, *no matter what you see*, learning about your future attitude does not tell Alice anything about who is in which cell, and so does not give her any information about which of the two cells contains the guilty suspect. Thus, there is no reason for Alice to change her confidence in Cell1 upon learning about your future, merely reasonable, attitude.

Here is the crucial point: since you know that Alice would be more confident in Cell 1 than Cell 2, *even upon learning about your future attitude*, you know that the evidence supports being more confident in Cell 1 than Cell 2. Thus, by DEFERENCE, you ought to be more confident in Cell 1 than Cell 2. But if we accept REASONABLE REFLECTION, it will follow that you should have no clue whether Cell 1 (since you

know that later you will have no clue whether Cell 1). So we should reject REASONABLE REFLECTION.

7.2 A worry about rejecting REASONABLE REFLECTION

In the previous subsection, I provided two reasons for thinking that we should reject REASONABLE REFLECTION. Once we reject REASONABLE REFLECTION, our problem is solved. We can have imprecise attitudes without facing a reflection violation because the correct version of reflection will not tell us to defer to our future attitude in the opaque detective case.

However, you might have the following worry about rejecting REASONABLE REFLECTION: Suppose I currently have some doxastic attitude towards p , but know that I will soon gain information which will make me reasonably adopt a different doxastic attitude towards p . Only rarely will I know that my later attitude towards p is best supported by the evidence (because, suppose, in most cases, my later attitude towards p will be imprecise). Nonetheless, it seems like, in many cases, I am better off deferring to my more informed attitude even if that attitude will not be the one supported by the evidence. So, if we deny REASONABLE REFLECTION, the reflection principle may not do the work we want it to do. Consider the following example:

MEDICAL TEST

You are a doctor trying to figure out whether your patient has disease D . If she does have D , it is crucial that you start treating her immediately. You have ordered some tests but have not seen the results yet. The technician then tells you that when you do look at the results, you will come to reasonably be very confident that the patient has D . Nonetheless, you know that you will not adopt a precise credence in the proposition that the patient has D .

It seems clear that you should now become confident that your patient has D and start treating the patient for D . This is true despite the fact that you will lack a future precise credence in that proposition. So why is it that, in some cases, we ought to defer to attitudes that are reasonable, yet not fully supported by the evidence, while in others we should not?

We can explain why you ought to defer to your future attitude in MEDICAL TEST without appealing to REASONABLE REFLECTION. In fact, all we need to explain why you should be confident that the patient has the disease is a generalized version of EVIDENTIAL REFLECTION:

GENERAL EVIDENTIAL REFLECTION: If you know that, in the future, you will have doxastic attitude A towards p , where A is *within the interval* R , and that the evidence will support a doxastic attitude towards p within R , then (provided you lose no information), you ought to now have a doxastic attitude towards p within R .

The same motivations for EVIDENTIAL REFLECTION motivate the generalized version (which was also defended by van Fraassen). How does this principle help? You may not know that you will have a future *credence* in the proposition that your patient has D that is supported by the evidence. You can, however, know that, in the future,

the evidence will support having a high degree of confidence (say, something between 0.8 and 1) in the proposition that your patient had D. After all, you know that this is a very reliable test for D. So GENERAL EVIDENTIAL REFLECTION dictates that you should now have a high degree of confidence that the patient has D.

Are there cases in which it is important to defer a future, merely reasonable, *credence*? I am not sure. But, even if there are, I do not think that REASONABLE REFLECTION would be the right way to explain these cases. This is because REASONABLE REFLECTION is much too simple a principle to guide us with regard to when we ought to defer to our future merely reasonable credences. As I mentioned earlier, in some cases, there are reasons to defer to our future judgment because that judgment is more informed than our current judgment, but also reasons to not defer to that judgment because that judgment is not as well supported by the evidence as our current one. The question, then, is going to be: What do we have *most* reason to do? Imagine that you have to consult with some experts about some matter. You know that one is more informed but the less informed one is better able to determine what is supported by the evidence. Which one will you trust?

There is no general answer to this question. Whether you defer to someone will depend on how likely you think they are to be right, and this will depend on how well informed they are *and* how well they respond to evidence. You could imagine a case where you defer to the more informed expert despite the fact that she is not quite as good at evaluating evidence. Alternatively, it might be that, although she has extra information, the additional information that she has is not enough to outweigh the consideration that the other expert's conclusion is better supported by the evidence. Or, perhaps, in some situations, you will not defer to either expert, but your degree of confidence will be somewhere in between their two attitudes. The point is this: Who you ought to defer to, and how much you ought to defer, is going to depend on very detailed information about the case at hand, and it is misguided to think that we can come up with a nice clean principle which will tell you, in every case in which there is one person who is better at evaluating evidence, and one who is more informed, which one you should defer to. Similarly, in cases in which your future judgment is more informed, but not as well supported by the evidence as your current one, we should not expect a principle which will tell us to always, or never, defer to that future attitude.²⁶

So far, all I have said is that whether or not you ought to defer to your later self is going to be a complicated issue. But what is clear is that there are some cases in which you definitely ought *not* defer to your later self, and these are the cases in which you know that the evidence supports maintaining your current attitude. For this reason, there is no need for us to worry about the fact that we lack precise credences. Recall that the worry for agents who violated PRECISION was that if an

²⁶ Briggs (Briggs 2009) has come up with a principle that is supposed to take these kinds of considerations into account. She has a reflection principle which she calls "Distorted Reflection." This principle tells you that if you know that your later credence in p will be r , and you would not lose any information between now and then, your credence in p now should be $r - D_r$, where D_r is a factor that expresses your expected departure from rationality. If we can formalize our expected departure from rationality (she has a suggestion as to how to do this as well) this may be exactly the kind of principle we need.

agent does not conform to PRECISION, there are cases in which she will not defer to her future attitude. But, in these cases, we know that she absolutely should not defer to her future attitude (since she knows that the evidence supports maintaining her current one). So, however the story of reflection principles ends, we can be sure that such principles will not cause problems for agents with imprecise attitudes.

8 Conclusion

In this paper, I presented a problem that arises when we think about how precise our doxastic attitudes need to be. I described arguments both for and against PRECISION, the claim that our doxastic attitudes need to be extremely precise. The argument against PRECISION was based on the idea that we are permitted to be insensitive to mild evidential sweetening, and the argument for PRECISION was based on the idea that, if we violate PRECISION, we might be forced to violate a plausible reflection principle. I have argued that we can solve this problem by appealing to the distinction between what we ought to believe and what the evidence supports. Once we recognize that the evidence can support precise credences, even though it is false that we ought have precise credences, we are problem free. For granting that the *evidence supports* precise credences means that we can still be insensitive to mild evidential sweetening, and granting that *we are permitted* to have imprecise attitudes will not yield a reflection violation.

One of the upshots of this is that when we come across very demanding principles of rationality, like PRECISION, we should stop and think about whether trying to use such a principle in deliberation will help us achieve our epistemic aims. If the principle does not have this feature, it is not a principle that we should use to deliberate with. However, this does not mean that such a principle has nothing going for it. For it might be that, even if having reasonable attitudes does not require meeting some very demanding condition C, having the attitudes supported by the evidence does. What I have argued for here is that, even in cases in which we should not have the attitudes that the evidence supports, understanding the evidential support relation is quite important. This is because, when we evaluate the epistemic credentials of another agent, or a future time slice of ourselves, the relevant question to ask is: How well are this agent's attitudes supported by the evidence? REFLECTION is a principle about deference, and so the right version of REFLECTION will say that, while we should always defer to future attitudes that are supported by the evidence, we should not always defer to our future reasonable attitudes.

I think that the distinction between what attitudes we ought to have and what attitudes the evidence supports is a powerful one, and that this distinction can be used to solve a variety of problems in epistemology. It is especially important that we recognize that the rules that govern *what we ought to believe* should be sensitive to our cognitive limitations, even though the rules that govern *what the evidence supports* should not account for such limitations. This is because, if we want the theory about what we ought to believe to help us reason, and if the purpose of reasoning is to approach the truth, the theory we adopt *must* account for our epistemic imperfections. After all, *we* are the ones that care about discovering the truth, and a theory that ignores our imperfections is not going to be as helpful as one

that accounts for them. Once we realize that the attitudes that we ought to have are not the ones that an agent with perfect cognitive capacities would have, we will, I hope, be better equipped to figure out what is true, and, as a plus, we can also enjoy some epistemic relaxation.

Acknowledgements In writing this paper, I have benefited greatly from conversations with David Christensen, Adam Elga, Daniel Greco, Caspar Hare, Eli Hirsch, Carrie Ichikawa Jenkins, Julia Markovits, Rebecca Millsop, Agustin Rayo, Susanna Rinard, Robert Stalnaker, Stephen Yablo and Roger White. I also received extremely helpful feedback from the audience at the Bellingham Summer Philosophy Conference, 2011, and members of the 2011–2012 MIT Job Market Seminar.

References

- Aarnio, M. L., (2010). “Unreasonable Knowledge.” *Philosophical Perspectives* 24, 1–21.
- Briggs, R. (2009). Distorted reflection. *Philosophical Review*, 118(1), 59–85.
- Chang, R. (1997). *Introduction to incommensurability, incomparability and practical reason*. Cambridge, MA: Harvard University Press.
- Christensen, D. (2008). Does murphy’s law apply in epistemology? Self-doubt and rational ideals. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 2). Oxford: Oxford University Press.
- Christensen, D. (2010). Higher order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- Earman, J. (1992). *Bayes or bust*. Cambridge, MA: MIT Press.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers’ Imprint*, 10(5), 1–11.
- Field, H. (2000). Apriority as an evaluative notion. In P. Boghossian & C. Peacocke (Eds.), *New essays on the a priori*. New York: Oxford.
- Hare, C. (2010). Take the Sugar. *Analysis*, 70(2), 237–247.
- Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach* (2nd ed.). Chicago: Open Court Publishing.
- Jeffrey, R. (1983). Bayesianism with a human face. In J. Earman (Ed.), *Testing scientific theories*. Minneapolis, MN: University of Minnesota Press.
- Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, 19, 153–178.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24, 281–323.
- Kaplan, M. (1996). *Decision theory as philosophy*. Cambridge: Cambridge University Press.
- Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy*, 71, 391–418.
- Levi, I. (1985). Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 52, 390–409.
- Sober, E. (2002). Bayesianism—Its Scope and Limits. In R. Swinburne (Ed.) *Bayes’ theorem* (vol. 113). Oxford: Oxford University Press.
- Strevens, M. (2005). The Bayesian approach in the philosophy of science. In D. M. Borchert (Ed.), *Encyclopedia of philosophy* (2nd ed.). Detroit: Macmillan Reference.
- van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81, 235–256.
- White, R. (2010). Evidential symmetry and mushy credence. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 3). Oxford: Oxford University Press.
- Schoenfield, Miriam (ms)^a. Why Acquaintance Matters”.
- Schoenfield, Miriam (ms)^b. Expecting too much of epistemic rationality: Why we need two notions instead of one.
- Schoenfield, Miriam (ms)^c. Permission to Believe.
- Sepielli, A (ms). Evidence, Reasonableness and Disagreement. “Unpublished”.