

Imprecision in Normative Domains

by

Miriam Schoenfield


B.A. Mathematics, Neuroscience, Philosophy
Brandeis University

SUBMITTED TO THE DEPARTMENT OF PHILOSOPHY IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN PHILOSOPHY
AT THE MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

© 2012 Miriam Schoenfield. All Rights Reserved



The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or part in any medium now known or hereafter created.

Signature of Author: _____

Department of Philosophy
May 10, 2012

Certified by: _____

Roger White
Professor of Philosophy
Thesis Supervisor

Accepted by: _____

Roger White
Professor of Philosophy
Chair of the Committee on Graduate Students

Imprecision in Normative Domains

by

Miriam Schoenfield

Submitted to the Department of Philosophy May 10, 2012 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Philosophy

ABSTRACT

Being rational and being moral can be difficult. However, some theories of rationality and morality make living up to these ideals *too* difficult by imposing requirements which are excessively rigid. In this dissertation, I defend and explore the implications of relaxing some of these requirements. I first consider the implications of thinking that rational agents' doxastic attitudes can be represented by imprecise, rather than precise probabilities. In defending this position, I develop a distinction between an idealized, and less idealized notion of rationality. I then explore the moral implications of the thought that facts about value cannot be represented by a precise value function. Finally, I defend *permissivism*, the view that sometimes there is more than one doxastic attitude that it is rationally permissible to adopt given a particular body of evidence, and show that this view has some interesting implications for questions about higher order evidence.

Thesis Supervisor: Roger White

Title: Professor of Philosophy

Table of Contents

Chapter 1: Chilling Out on Epistemic Rationality – A defense of imprecise credences and other imprecise doxastic attitudes.....	4
Chapter 2: Why Acquaintance Matters – imprecision in value and the ethics of helping.....	33
Chapter 3 Permission to Believe – why permissivism is true and what it tells us about irrelevant influences.....	57

Chapter 1

Chilling Out on Epistemic Rationality¹

A Defense of Imprecise Credences (And Other Imprecise Doxastic Attitudes)

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, our belief-like states should be entirely representable by a precise probability function

¹ In writing this dissertation, I have benefited greatly from the help of my dissertation committee, supervised by Roger White and including David Christensen, Caspar Hare, Julia Markovits and Robert Stalnaker. I am also thankful for helpful conversations with Tyler Doggett, Tom Dougherty, Adam Elga, Daniel Greco, Aaron Hauptman, Eli Hirsch, Sophie Horowitz, Carrie Ichikawa Jenkins, Rebecca Millsop, Agustin Rayo, Susanna Rinard, Yosef Schoenfield, Stephen Yablo, and Palle Yourgrau. Finally, I received extremely helpful feedback from the audience at the Bellingham Summer Philosophy Conference, 2011, and members of the 2011-2012 MIT Job Market Seminar.

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 section two, 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 sections three and four, 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 section five, 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 .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, twenty four 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.² 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.³ For this reason, there have been attempts to construct a version of Bayesianism that can survive the rejection of PRECISION. (As Isaac Levi so poetically notes, those who "object to the numerical precision required by strict Bayesianism are accommodated in the ample bosom of Mother Bayes").⁴ 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.

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 people⁵ 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.

² See, for example, Sober (2002).

³ See, for example, Earman (1992), Howson and Urbach (1993) and Strevens (2006).

⁴ Levi (1985, 392).

⁵ See, for example, Levi (1974), Jeffrey (1983), Kaplan (1996), Joyce (2005) and Joyce (2010).

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 $\sim p$.^{6,7}

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)

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?

⁶ 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.

⁷ For discussions of insensitivity to sweetening in the context of practical rationality and ethics see, for example, Chang (1997), Hare (2010), and Schoenfield (ms.)⁸.

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 crimes, "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*.⁸ 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).

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.⁹ 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 .¹⁰

⁸ For early discussions of this model see Jefferey (1983) and Levi (1985).

⁹ In fact, there are a number of compelling arguments for PRECISION (see Elga (2010) and 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.

¹⁰ 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.

This principle is in the same spirit as the Reflection principle that was first introduced and defended by Bas 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 *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 Cell #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 in to 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 “Cell1” and the proposition that the person in Cell #2 committed the crime “Cell2.” 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 Cell1 than in Cell2.

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 Cell1 or Cell2.
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 Cell1 than Cell2 when you learn about the additional eyewitness. I will present three arguments for this claim:

Argument #1: The Reasons for Belief Argument

What reasons do you have to believe Cell2? 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 Cell1? 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 Cell2 are also reasons to believe Cell1. However, you have an *additional* reason to believe Cell1 which is not a reason to believe Cell2, namely that the eyewitness claims to have seen the person in Cell #1 commit the crime. Since every reason you have to believe Cell2 is also a reason to believe Cell1, but you have an additional reason to believe Cell1 that is not a reason to believe Cell2, it follows that you have more reason to believe Cell1 than Cell2.¹¹

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 Cell1 than Cell2. To see

¹¹ This argument was inspired by a similar argument in decision theory described in Hare (2010).

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 Cell1 and 0.5 credence to Cell2.¹² 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 Cell1 than Cell2 (and thus *you* are more confident in Cell1 than Cell 2).

Argument #3: The Argument from $Pr(heads) = 0.5$

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.

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

Recall that Cell1 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 Cell1 is true if and only if H is true. Since it is clear that you should have a 0.5 credence in H,

¹² 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) Cell 1 \leftrightarrow ((S&H) or (J&T)) (in other words, the person in Cell1 is guilty if and only if Smith is guilty and the coin landed heads or Jones is guilty and the coin landed tails).
- (2) $Pr(\text{Cell1}) = Pr((S\&H) \text{ or } (J\&T))$
- (3) $Pr(\text{Cell1}) = .5Pr(S) + .5Pr(J)$
- (4) $Pr(J) = 1 - Pr(S)$ (since either Smith or Jones is guilty)
- (5) $Pr(\text{Cell1}) = .5Pr(S) + .5(1 - Pr(S)) = 0.5$
- (6) $Pr(\text{Cell2}) = 1 - Pr(\text{Cell1}) = 0.5$

and since your attitude towards H and Cell1 must be the same, you should have a 0.5 credence in Cell1 as well.¹³ But if you have a precise credence in Cell1, you will be sensitive to mild evidential sweetening, and so, when the eyewitness comes in, you will be more confident in Cell1 than Cell2. Thus, at very least, in this version of the opaque case, you should be more confident in Cell1 than Cell2.¹⁴ 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 Cell1 or Cell2. 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 Cell1 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 Cell1 than Cell2 and no more confident in Cell2 than Cell1.

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 Cell1 than Cell2 and no more confident in Cell2 than Cell1. So you know that, *no matter how things turn out*, you will later have no clue whether Cell1 or Cell2. But then it seems that, according to REFLECTION, you ought to not have a clue whether Cell1 or Cell2 *even before you find out who is in which cell*, contrary to the judgment

¹³ There is an argument in White (2010 p.175-181), which, in this case, could be applied to defend the claim that you should match your credence in Cell1 to your credence in H, rather than the other way around.

¹⁴ 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).

that, in the opaque case, you ought to be more confident in Cell1 than Cell2.¹⁵ In the next 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.¹⁶ I defend this distinction in greater detail elsewhere¹⁷ 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*.¹⁸ 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

¹⁵ 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 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.

¹⁶ 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 David Christensen (2008), and is necessary for agent neutrality, which will be discussed shortly. (This kind of perfect rationality is related to the notion Harry 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.

¹⁷ Schoenfield (ms.)^b. Also, see Aarnio (2010) and Sepielli (ms.) for discussions of distinctions along these general lines.

¹⁸ 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*.

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.¹⁹

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 pi 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 shouldn’t be confident in it now, despite the fact that the experts say we should. Thus, it can’t be right that we should be confident in all of the mathematical truths, and so it can’t 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,²⁰ but has no reason to think that *she*

¹⁹ 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 doesn’t depend on which *particular* agent is evaluating it (though what the evidence supports will depend on the agent’s priors, or standards of reasoning).

²⁰ 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.

is under the influence of a drug, should not reduce her confidence in *p*.²¹ Since I am assuming that what the evidence supports does not depend on which particular agent is evaluating it, the evidence can't 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. Maybe. 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 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

²¹ This phenomenon was first discussed (as far as I know) in Christensen (2010).

you prefer, maximizing her expected accuracy²²).²³ 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 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.

²² 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.

²³ 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.

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 section seven). This is how the competing pressures concerning precise credences will get resolved. I will spell out the solution in more detail in section seven, 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* shouldn't 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 *can't* always support a precise credence. (If this is not a concern of yours, feel free to skip to section seven, in which I spell out the solution). James Joyce (2005, 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, 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,²⁴ 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

²⁴ 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.

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 intervals, these intervals would be varied.²⁵ Perhaps the fully rational agent would have credence [.2-.3] towards the proposition about elephants, but [.4-.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).²⁶

²⁵ 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 14).

²⁶ I argue against the *uniqueness* assumption in Schoenfield (ms)^c.

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 *rationally* have doxastic attitude A towards p, without any loss of information, you ought to now have doxastic attitude A towards p.

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 Cell1 or Cell2, 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 Cell1, 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 Cell1 even before you learn who is in which cell, contrary to our assumption. Fortunately, there are good reasons to reject REASONABLE REFLECTION.

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.

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 Cell1 than Cell2 is the attitude supported by the evidence, and yet, REASONABLE REFLECTION would say that you should have no clue whether Cell1 or Cell2. 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 Cell1, 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 Cell1 would depend on who turns out to be in Cell1. 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 Cell1 than Cell2. 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 Cell1. Would this affect *her* confidence in Cell1?

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 Cell1, *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 Cell1 than Cell2, *even upon learning about your future attitude*, you know that the evidence supports being more confident in Cell1 than Cell2. Thus, by DEFERENCE, you ought to be more confident in Cell1 than Cell2. But if we accept REASONABLE REFLECTION, it will follow that you should have no clue whether Cell1 (since you know that later you will have no clue whether Cell1). 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 .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.²⁷

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 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

²⁷ Rachael 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 won't 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.

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.

References

- Aarnio, Maria Lasonen (2010). "Unreasonable Knowledge." *Philosophical Perspectives* 24
- Briggs, Rachael (2009). "Distorted Reflection." *Philosophical Review* 118 (1)
- Chang, Ruth (1997) Introduction to *Incommensurability, Incomparability and Practical Reason* Harvard University Press: Cambridge, MA.
- Christensen, David (2008). "Does Murphy's Law Apply in Epistemology? Self-doubt and Rational Ideals." *Oxford Studies in Epistemology Vol 2*. Tamar Szabo Gendler and John Hawthorne eds. Oxford University Press. Oxford, UK.
- Christensen, David (2010). "Higher Order Evidence." *Philosophy and Phenomenological Research*. 81(1).
- Earman, John (1992). *Bayes or Bust*. MIT Press. Cambridge, MA.
- Elga, Adam (2010). "Subjective Probabilities Should be Sharp." *Philosophers' Imprint* 10 (5).
- Howson, Collin and Urbach, Peter (1993). *Scientific Reasoning: The Bayesian Approach*. 2nd Edition. Open Court Publishing. Chicago, IL.
- Hare, Caspar (2010). "Take the Sugar." *Analysis* 70 (2).
- Jeffrey, Richard (1983). "Bayesianism with a Human Face" in *Testing Scientific Theories*. John Earman ed. University of Minnesota Press. Minneapolis, MN.
- Joyce, James M (2005). "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19.
- Joyce, James M (2010). "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24.
- Kaplan, Mark (1996). *Decision Theory as Philosophy*. Cambridge University Press. Cambridge, UK.
- Levi, Isaac (1974). "On Indeterminate Probabilities." *Journal of Philosophy* 71.
- Levi, Isaac (1985). "Imprecision and Indeterminacy in Probability Judgment," *Philosophy of Science* 52.
- 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, Andrew (ms). "Evidence, Reasonableness and Disagreement." Unpublished manuscript.

van Fraassen, Bas (1984). "Belief and the Will." *Journal of Philosophy* 81: 235–56.

Sober, Elliott (2002). "Bayesianism - Its Scope and Limits" in *Bayes' Theorem*. Richard Swinburne, ed. Proceedings of the British Academy Press, vol. 113.

Strevens, Michael (2005). "The Bayesian Approach in the Philosophy of Science" in *Encyclopedia of Philosophy, second edition*. D.M. Borchert ed. Macmillan Reference.

White, Roger (2010). "Evidential Symmetry and Mushy Credence." *Oxford Studies in Epistemology Vol. 3*. Tamar Szabo Gendler and John Hawthorne eds. Oxford University Press. Oxford, UK.

Chapter 2

Why Acquaintance Matters *Imprecision in Values and the Ethics of Helping*

1. Introduction

You don't have to be a consequentialist²⁸ to care about value. Almost everyone will agree that *sometimes* what you ought to do is governed, at least *in part*, by the value, or expected value²⁹, of our actions. For example, suppose you have a choice between curing Alice's headache or Bob's cancer. All else equal, you should cure Bob's cancer. Why? Seemingly because Bob's having cancer is *worse* than Alice's having a headache. This example involves comparative judgments about the badness of harms that befall different people (interpersonal comparisons of value) and some people are suspicious of such comparisons. Nonetheless, such a person might still want to appeal to value judgments to explain why it is that if you have a choice between curing *Alice's* headache or *Alice's* cancer, you should cure Alice's cancer. These judgments motivate what I will call "minimal consequentialism."

MINIMAL CONSEQUENTIALISM: Suppose you are faced with a choice between two ways of helping: doing A or doing B. And suppose further that neither choice will involve you doing something wrong in itself,³⁰ and that you have no personal preference or associative obligation³¹ to the people involved that would dictate which of A or B you should do. Then, the permissibility of doing A and doing B should be determined by the comparative expected value of doing A and doing B.³²

Although MINIMAL CONSEQUENTIALISM seems quite plausible, I am going to argue that, if we accept MINIMAL CONSEQUENTIALISM, we will face a problem that results from imprecision

²⁸ Consequentialism is the view that what matters morally is the consequences of our action. In particular, many think that what we ought to do is completely determined by how *good* or *valuable* the outcome of an action is.

²⁹ I will talk in greater detail about expected value in section three. For now, it suffices to think of the expected value of some action as a way of measuring the action's "worth" when we don't know exactly what will come about as a result of the action.

³⁰ For example, breaking a promise, violating someone's rights, or lying.

³¹ For example, obligations arising from friendship, or familial relationships.

³² If you don't like interpersonal comparisons of value, you will prefer the following version of the claim: Suppose you are faced with a choice between two ways of helping *S*: doing A or doing B.... Then, the permissibility of doing A and doing B should be determined by the comparative expected value *for S* of doing A and doing B. This difference will not matter for what follows.

in values.³³ I will argue that to solve the problem we need to think that, when choosing how to help, even when standard non-consequentialist constraints are satisfied, there is more to consider than the values of the outcomes. There is another very important, and often overlooked, consideration that is relevant to decisions about how to help, and this is the degree to which we are *acquainted* with the reasons to help. What is it to be well acquainted with the reasons to help? Roughly, it is to have a lot of detailed information about the reasons to help. For example, suppose you know that Carrie needs a hundred dollars to buy some medicine. This gives you a reason to help Carrie. Then suppose you learn that Carrie needs the one hundred dollars to buy some *asthma* medication. Learning that the medication is for asthma does not provide you with an *additional* reason to help Carrie. It does, however, mean that you are better *acquainted* with the reasons to help Carrie. I want to suggest that facts about acquaintance are morally relevant. All else equal, the greater our acquaintance with the reasons to help someone, the greater our obligation to help is.

Here is the plan: in section two, I am going to talk about imprecision in value and how it plays a role in our decisions of whom to help. In section three, I will describe a problem that arises as a result of imprecise values. In section four, I will suggest a possible solution which appeals to facts about the metaphysics of value. Ultimately, however, I will suggest that this solution is unsatisfactory and so, in section five, I will propose that we can solve the problem by appealing to the moral relevance of acquaintance. I will conclude by showing that recognizing the moral relevance of acquaintance can help explain a number of puzzling issues in ethics.

2. Imprecise Values

In this section, I will explore a phenomenon that arises when we allow for imprecision in value called *insensitivity to sweetening*. This phenomenon will provide the basis for the problem that, I think, arises when we take value facts to be imprecise. In section 2.1, I will describe insensitivity to sweetening. In section 2.2, I will suggest an explanation for why we might be insensitive to sweetening and, in section 2.3, I will describe how we might model this phenomenon for the purposes of decision making under conditions of uncertainty.

³³ A version of this problem will arise for those who accept the version of MINIMAL CONSEQUENTIALISM described in the previous footnote.

2.1 Insensitivity To Sweetening

As an example of the phenomenon that I am interested in, I will describe a case in which, seemingly, how to help should be determined by the values of the possible outcomes. The case is one in which you are choosing between donating your money to two different charities, each serving a different population. However, if you do not like interpersonal comparisons of value, you can just as well imagine a case in which the two charities serve the same population, but in different ways.

CHARITY TRANSPARENT

You have come into a large sum of money and want to donate it to charity. You have two options. Either you could expand the free meal program for the homeless to include an extra 50 meals a day, or, you could start a program in a low-income school which would enable 500 students a year to study some form of art (painting, dance, music, theatre, etcetera.).

You have no associative obligations or personal preference about which organization to help and neither option involves performing an action that is wrong in itself. Let us assume for now that MINIMAL CONSEQUENTIALISM is true and that, therefore, the only relevant consideration, in this case, is how much good you will do by giving to either cause. Which of these causes should you support?

Suppose that you have done a great deal of research on the impacts of programs like these and you have a lot of incredibly detailed information about the programs. Here are some examples of some of the conclusions you have drawn: If you support the arts program, over the years, approximately 1.5% of students will be inspired to pursue art and will make valuable contributions to the arts community. If you support the meal program, there will be people who, as a result of the savings on food, will be able to afford medication that is very important to their health. The arts program will employ some struggling artist to teach the students. The meal program will provide much needed additional nutrients to pregnant women who are at high risk. The arts program will keep kids engaged in a productive activity and some studies show that such activities will decrease their chances of getting involved in drugs and violence. The expansion of the meal program will mean that the venue in which the meals are served is open for two more hours a day and can therefore provide more hours of shelter during bad weather. The list goes on.

You know a lot about what will happen as a result of donating to each charity and you have thought hard about all of the considerations. You reasonably conclude that the expected value of donating to one is no greater or less than the expected value of donating to the other.³⁴ Note that this judgment is not necessarily a result of ignorance about how things will go. We can design the case so that (albeit unrealistically) you know all the facts about how the world would go were you to donate to one versus the other program and yet you would *still* reasonably judge that the value of donating to one is no greater or less than the value of donating to the other.

Now suppose that, after all of this deliberation you learn that, actually, the school has 501 students, rather than 500. This means that if you donated money to the school, 501, rather than 500 students a year would benefit from the program that you would start. Upon learning this information, are you now obligated to donate to the school? It seems not. If it was permissible to give to either cause before, it seems like it is permissible to give to either cause now.³⁵ And note that *any* way of slightly increasing the value of one outcome over the other would not be enough to break the tie. If, for example, by donating to one rather than another you could relieve a stranger's mild headache, obtain an extra ten dollars to give to your chosen organization, plant some pretty flowers in a city garden or enable the person on the bus next to you to have a better hair day – it still would not be the case that you should donate to one rather than the other. In cases in which small differences do not break the tie, I will say that you are *insensitive to sweetening*.³⁶

Related to the notion of insensitivity to sweetening is the notion of two outcomes being *on a par*.³⁷ Outcomes A and B are on a par if and only if there exists an outcome, A+, such that the following is the case:

- a) A is no better than B and B is no better than A.
- b) A+ is better than A.
- c) A+ is no better than B and B is no better than A+.

In the case we just discussed, A is the outcome in which you give your money to a school with 500 students, A+ is the outcome in which you give your money to a school with 501 students and

³⁴ It is not crucial to the argument that you agree with this judgment given the particular way I spelled out the case. If you do not agree with the judgment, just fill in the details however you please so as to get the result that the expected value of one is no greater or less than the value of the other.

³⁵ I should note that a classical hedonic utilitarian might disagree with this judgment.

³⁶ For a discussion of insensitivity to sweetening in other contexts (like practical and epistemic rationality) see Hare (2010) and Schoenfield (forthcoming).

³⁷ Here, and in much of my discussion, I follow Ruth Chang's (1997) terminology.

B is the outcome in which you give your money to the meal program. While A+ is better than A (note that if you had to choose between the program being available for 500 students, or the program being available for 501 students, clearly you should make the program available to 501 students!), A is no better than B, and B is no better than A. Furthermore, since small differences do not break the tie, A+ is no better than B and B is no better than A+. Thus, in CHARITY TRANSPARENT you are facing a choice between two outcomes which are on a par.

If there exist outcomes which are on a par, the facts about value cannot be represented by a precise value function which assigns to each outcome a unique real number, representing how valuable the outcome is. The reason is simple. 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 $r_1=r_2$, and if $r_1=r_2$, then any number greater than r_1 will also be greater than r_2 . Let $V(x)$ denote the value of x . If we start out assuming that A is no better than B and B is no better than A, and that the values of A and B can be represented by real numbers, then $V(A) = V(B)$, and if $V(A+)$ is a number greater than $V(A)$ it will, therefore, also be greater than $V(B)$.

2.2 Why We Are Insensitive to Sweetening

Here is one account of what is going on in cases in which we are insensitive to sweetening: When I am looking at the two options, I am trying to decide which option would bring about the most good *overall*. But there are lots of factors that contribute to how good an outcome is. Health, beauty, justice, knowledge, happiness, and pleasure are just some examples of factors that might contribute to the goodness of an outcome. Let us call these values the *contributory values* and the value of the outcome as a whole *the covering value*. The problem we face is this: For most of us, there is no way of weighing the contributory values which seems to us to be the uniquely correct one. If there were, we would be able to answer questions like this: Take outcomes A and B and suppose that outcome B is *marginally* better than outcome A. *Exactly* how many more minutes of sunshine would we have to add to A such that if we added a *minute* more, the balance would be tipped and outcome A would be better? Or, exactly how many delicious chocolate bars is the creation of a piece of mediocre art worth? These are questions that we have no answers for, and our inability to assign values in this precise way leads to our judging outcomes to be on a par. In the charity case, there are several contributory values and the problem we face is that we lack a precise way of weighing these values against each

other. If, for values V_1 and V_2 , there is no way of weighing these values which seems to S to be the uniquely correct one, I will say that these values are *incommensurable* for S .³⁸ Note that the way I am using the term “incommensurability” allows that a subject can face two options with contributory values that are incommensurable for her, and yet one option might be definitively better than the other. Just because there is no *unique* way of weighing the two values does not imply that *all* ways of weighing them are acceptable. For example, if the choice was between an arts program and a program which would solve world hunger permanently, the latter program would bring about greater value than the former, even though, we may have no unique way of precisely weighing the value of art against the disvalue of hunger.³⁹

To further clarify this thought, it might be helpful to contrast this case with one in which two contributory values are most definitely *not* incommensurable for us. Take, for example, the covering value to be wealth, and the contributory values in question to be the number of dollar bills and the number of quarters. In this case, there *is* a single way of weighing the contribution of the contributory values which seems to us to be uniquely correct: four quarters contribute the same amount to the covering value (wealth) as does one dollar bill. Note that in the case of the quarters and dollar bills, we are never insensitive to sweetening. Given two sums of money composed of a mixture of quarters and bills, if we know that neither sum is of greater value than the other, this will be because they are each worth *exactly* the same amount of money, and adding a penny to one *would* break the tie.

2.3. Imprecision and Acting Under Conditions of Uncertainty

So far, I have described insensitivity to sweetening and suggested that incommensurability in values might explain this phenomenon. In this subsection, I will be dealing with the question of how insensitivity to sweetening affects the standard models used to

³⁸ I should note that I use the term “incommensurability” in a slightly different way than Ruth Chang (1997) does.

³⁹ I should note that other explanations have been offered for insensitivity to sweetening in moral cases. Francis Kamm (2000) and Tyler Doggett (2009) discuss cases in which you should not be swayed by a small sweetening and their explanation of this phenomenon appeals to the claim that being swayed by a sweetening would be *disrespectful*, as it suggests that you are not taking the needs of the people involved sufficiently seriously. I have a number of worries about this approach: first, it will only work in an *interpersonal* case. If I am choosing between two ways of helping the same person, it cannot be disrespectful to the person being helped that I’m swayed by a sweetening (especially if the sweetening, itself, is a benefit to the person) to help in one way rather than another. Second, it seems like we should be insensitive to sweetening even in situations in which the needs are not particularly serious, and so an explanation that appeals to the demand that we take the needs in question quite seriously may not be sufficiently general.

represent decision making under conditions of uncertainty. (I will be restricting my discussion here to cases in which we are only concerned with the values of the outcomes).

If we are not omniscient (which, I'm afraid, is in fact the case), we do not always know what the outcome of our actions will be. For this reason, many consequentialists appeal to expected value theory to give an account of what you ought to do under conditions of uncertainty. The idea is as follows: Even in situations in which you do not know the value of choosing A and the value of choosing B, you can still know the *expected value* of choosing A and B respectively. What is the expected value of a choice? It is the sum of the values of the possible outcomes weighted by the probability of those outcomes occurring. If we do not have a precise value function (in other words, if there are outcomes that are incommensurable for us) expected value theory might be in trouble. For calculating the expected value of some choice requires plugging in *numbers*. If we don't have numbers, expected value theory will not work.

Fortunately, things are not as dire as they may seem. What we can do, if we do not have a single value function, is consider a set of value functions. Since we have a set of value functions, rather than a single function, the value of some outcome will not be represented by a number, but rather, by an interval. Let us call the model, which uses sets of functions, instead of one unique function, *the supervaluationist model*. Each function in the set will represent a different (acceptable)⁴⁰ way of weighing the different values involved. We can use this set of functions to compare the values of two outcomes as follows: Once again, let $V(x)$ denote the value of X.

- 1) For any outcomes A and B, $V(A)$ is greater than $V(B)$ if and only if every function U in the set assigns a higher value to A than to B.
- 2) For any outcomes A and B, $V(A)$ is equal to $V(B)$ if and only if every function U in the set assigns equal value to A and to B.
- 3) For any outcomes A and B, A and B are on a par if and only if there exists a function U in the set which assigns a higher value to A than to B, and another function U' which assigns to A a lesser or equal value than it does to B.

How do we calculate expected value to determine what we ought to do when we have a set of value functions instead of just a single function? Suppose that we have to decide whether

⁴⁰ When I say the value function is "acceptable" I don't mean that the value function would be an acceptable way of representing the facts about value (for you might think that no single value function can successfully represent the facts about value). Rather I mean that the value function provides an acceptable precisification.

we should perform act X or act Y. We first calculate the expected value of doing X according to each of the individual functions in the set. We then calculate the expected value of doing Y according to each of the individual functions in the set. The supervaluationist suggestion is that it is impermissible to do X if and only if, according to each function in the set, the expected value of doing Y is greater than the expected value of doing X.⁴¹ This will be important later on.

3. The Problem

The goal of this section is to present a problem that arises when we are insensitive to sweetening. Here is the case we will be considering:⁴²

CHARITY OPAQUE

As before, you have a choice between donating to a meals program or an arts program. But this time, you are in a rural area and you need to deliver the money in cash to the office of the program that you decide to donate to. You've been given two addresses, one for each office. They are:

50 East St

50 West St

You are weighing up all of the considerations in favor of each option when you suddenly realize that you forgot to write down which address corresponds to which charity. You are in the middle of nowhere, with no phone or internet access, and you only have a couple of hours before you need to start traveling to the airport. "Well," you think to yourself "so much for all these complicated considerations, I guess I'll just choose which direction to go randomly." As you are getting into your jeep, still undecided about where you are going, you see in your rear view mirror a man with a large suitcase, heading eastward. You think to yourself "Hmm...if I went east I could help the weary traveler out by giving him a ride."

⁴¹ The supervaluationist version of expected value theory can be challenged as follows: You might think that it is just as implausible to suppose that we can represent the value of some outcome by an interval with sharp boundaries on the real number line as it suppose that the value of an outcome can be represented by a real number. This is the problem of higher order vagueness that supervaluationists about anything run into. I do not pretend to have a solution to this problem, but it seems like there are two routes worth pursuing. First, one might allow the boundaries themselves to be vague. David Christensen (2004, p.150), for example, suggests that this might be a promising solution. Second, one might think that it is less problematic to have values represented by intervals with sharp boundaries than it is to have values represented by real numbers. John Broome (1997) has an argument which suggests that this might be the case. I should also note that there is an alternative model, described in Hare (2010), which yields different results in the cases that I will be discussing. However, for reasons I cannot get into here, I think the supervaluationist model is preferable.

⁴² Once again, if you worry about interpersonal value judgments, just imagine that the case involves choosing between two ways of helping the same population. The sweetening can also involve a benefit to that same population.

Question: what should you do? It seems obvious. You should go east and give the weary traveler a ride. If this does not seem obvious to you, consider the following two arguments for this claim which I will call, for convenience, HELP TRAVELER.

HELP TRAVELER: In CHARITY-OPAQUE you should help the traveler.

Argument #1 for HELP TRAVELER: The Supervaluationist Argument

One way to support the claim that in CHARITY OPAQUE you should go east is to appeal to the formal apparatus used to model insensitivity to sweetening. Let's consider the set of value functions that you are using in determining which outcome has the greatest expected value. Although the set of functions includes some functions which assign a higher value to the meal program than to the arts program and other functions which assign a higher value to the arts program than to the meal program, all of the functions will agree that going east has a greater expected value than going west.⁴³ According to the supervaluationist model, if all of your functions assign a higher expected value to going east than they do to going west, then going east has greater expected value than going west. If value is all we are concerned with in this case, it will follow that you should go east.

Argument #2 for HELP TRAVELER: The Reasons Argument

Another reason why you ought to go to east is that it would be quite insensitive, when you know nothing at all about which charity is in which location, to completely ignore the fact that by going east, rather than west, you could help the weary traveler. To see this, note that every reason you have to go west, is also a reason to go east. Since you don't know which

⁴³ Take any function U and let us suppose that U assigns the following values:

O_1 : You donate to the meal program and don't give the traveler a ride. Value of $O_1 = x$

O_{1+} : You donate to the meal program and give the traveler a ride. Value of $O_{1+} = x+n$

O_2 : You donate to the arts program and don't give the traveler a ride. Value of $O_2 = y$

O_{2+} : You donate to the arts program and give the traveler a ride. Value of $O_{2+} = y+n$.

Assume that your credence that the meal program is East = your credence that the arts program is East = 0.5. The expected value of going east is:

$$\Pr(O_1/\text{You go east})V(O_1) + \Pr(O_{1+}/\text{You go east})V(O_{1+}) + \Pr(O_2/\text{You go east})V(O_2) + \Pr(O_{2+}/\text{You go east})V(O_{2+}) = (0)(x) + (0.5)(x+n) + (0)(y) + (0.5)(y+n) = 0.5x+0.5y + n$$

The expected value of going west is:

$$\Pr(O_1/\text{You go west})V(O_1) + \Pr(O_{1+}/\text{You go west})V(O_{1+}) + \Pr(O_2/\text{You go west})V(O_2) + \Pr(O_{2+}/\text{You go west})V(O_{2+}) = (0.5)(x) + (0)(x+n) + (0.5)(y) + (0)(y+n) = 0.5x+0.5y.$$

Since $0.5x+0.5y+n > 0.5x + 0.5y$ the expected value of going east is greater than the expected value of going to west according to U. Since U was any arbitrary function, we can conclude that it is true of all functions that the expected value of going east is greater than the expected value of going west.

program is in which location, your reason to go west is just the fact that by going west you will help *either* the meals program or the arts program, and this is also a reason to go east. Furthermore, you have an *additional* reason to go east which is not a reason to go west, namely, that by doing so you will be able to help the weary traveler. Since all of your reasons to go west are also reasons to go east, but you have an additional reason to go east, which is not a reason to go west, it follows that you have *most* reason to go east.⁴⁴ (Note that this argument requires thinking of your moral reasons as *subjective* – in other words, it requires thinking what moral reasons you have depends on your epistemic situation. This is a commonly, though not universally, accepted way of understanding moral reasons. See, for example, Graham (2010) for an exception).

In fact, if you fail to go east, it seems as if you are taking the traveler's discomfort to provide you with *no reason at all*. For if it provided you with any reason, you would judge that you have most reason to go east. This is because all of the other reasons for going east and west are exactly the same. But why shouldn't the traveler's discomfort provide you with *any* reason to go east? After all, the traveler is a human being, who is tired, hot and uncomfortable!

Note that this argument could not be given if you knew which charity was at which address. For then it would not be the case that every reason you have to go west is a reason to go east. If you knew that the arts program was on West St, then you would know that by going west you can give students an opportunity for artistic expression, and there are strong reasons in favor of you doing that, which are not reasons to go east. Thus, if you failed to help the traveler, it wouldn't be because you were failing to see his discomfort as providing a reason. You would just think that the reason wasn't strong enough to outweigh all of the considerations to go west.

We now have some reasons to think that in CHARITY OPAQUE you ought to go to east: the first appealed to the supervaluationist model, the second appeals to the importance of seeing the traveler's discomfort as providing a reason. The problem with this judgment is that, once you find out which charity is in which location, you will not regard the outcome you brought about as any better than the alternative. To see this, imagine that you found out that the *meals* program is on East St. Since you are insensitive to sweetening, you do not regard the outcome in which you support the meals program and help the traveler to be any better than the outcome in which you support the arts program. Now imagine that you found out that the *arts* program is on East St.

⁴⁴ This argument in favor of going east is inspired by a similar argument in Hare (2010).

Since you do not regard the outcome in which you support the arts program and help the traveler to be better than the outcome in which you help the meals program, you still will not think that the outcome you brought about is better than the alternative would have been. In other words, *however things turn out*, you will not judge the value of the resulting outcome to be any better than the alternative. The worry, then, is a sort of reflection worry. How can it be the case that you ought to go east, while knowing that, once you get more information, you will rationally judge that neither outcome is better than the other (recall that we are supposing, for now that, in this situation, the value of the outcome is the only consideration)?

I am going to make the worry more precise by appealing to a principle that I will call CHOICE REFLECTION.

CHOICE REFLECTION: Suppose that you are faced with a choice between doing X and doing Y. If you know that, at a later time, you will rationally, without loss of information, judge that $EV(X)$ bears a certain relation R (greater than, less than, no greater and no less than) to $EV(Y)$, you should now judge that $EV(X)$ bears R to $EV(Y)$.

The simplest (and most persuasive, I think) reason for accepting CHOICE REFLECTION is that your future self's expected value judgments are based on more information, and so, you should treat your future self as an expert. That is why you should defer to her judgments. Another way to motivate CHOICE REFLECTION is to show that, if you violate CHOICE REFLECTION, you are subject to a diachronic Dutch Book; in other words, there is a series of actions that would each, individually, be warranted by your value judgments which will guarantee a sure loss.⁴⁵

As long as we assume MINIMAL CONSEQUENTIALISM (and thus that the only relevant consideration in this case is the value of the outcome) then, the judgment that you ought to help the traveler is inconsistent with CHOICE REFLECTION. The problem, then, is as follows: We have strong reasons to think that you *should* help the traveler, but also an argument for the claim that it's permissible *not* to help the traveler. This is not a happy situation. In the next section, I will

⁴⁵ Let's consider two cases. In the first case you know that later you'll judge $V(A) > V(B)$ while you now judge that $V(A)$ is not greater than $V(B)$. Suppose you start with A. Since you judge $V(A)$ to be no greater than $V(B)$, you will be willing to exchange A for B. But once you judge that $V(A) > V(B)$ you'll be willing to pay some amount of money, m to get A back. So you will have lost m. In the second case you know that later, you will judge $V(A)$ to be no greater or less than $V(B)$, while you now judge $V(B)$ to be greater than $V(A)$. Suppose you start with A. You'll be willing to pay m to exchange A for B. But once you judge that $V(A)$ is no greater or less than $V(B)$ you'll be willing to exchange B back for A. You'll have lost m. In both case, you have all the information that the bookie has (since you know what your future attitudes will be).

propose one solution, arising from metaphysical considerations, but ultimately, I will argue, that it is unsatisfactory. To successfully solve the problem, I will argue, we need to give up MINIMAL CONSEQUENTIALISM and appeal to a different kind of moral consideration, and that is the degree to which are acquainted with the reasons to help.

4. Metaphysics to the Rescue?

The problem in the previous section arose because, although seemingly we should help the traveler in the opaque charity case, this judgment is inconsistent with CHOICE REFLECTION (the principle that says to defer to your future, more informed judgments about relative expected value). In this section I will argue that we may be able to solve the problem by an appeal to metaphysics. However, ultimately, I do not find the metaphysical solution satisfying because it relies on some rather hefty metaphysical commitments. In section five, I will show that we can avoid the problem entirely and get everything we want by rejecting MINIMAL CONSEQUENTIALISM. In fact, this is what I think we should do. But let us first examine how metaphysics can potentially save the day.

Here are two options for how the metaphysics involved in mushy values might go:

MUSHY WORLD: There really are outcomes A, A+, and B such that A+ is better than A, A is no better than B, B is no better than A, but A+ is also no better than B. (In other words, there really are outcomes which are on a par).^{46,47}

MUSHY HEAD: There are no outcomes which are on a par. Rather our insensitivity to sweetening is a result of our ignorance of the correct way to measure value.

According to MUSHY WORLD, the incommensurability that we experience stems from the nature of (objective) value itself. Our insensitivity to sweetening reflects the fact that value, in

⁴⁶ There are two versions of the MUSHY WORLD view which I won't be distinguishing between. According to the first version (advocated by, for example, Raz (1985) and Chang (1997)), there exist outcomes A, A+ and B such that it is false that $V(A)$ is greater than $V(B)$, false that $V(B)$ is greater than $V(A)$ and false that they are equally valuable. According to the second version (advocated by Broome (1997) and Constantinescu (forthcoming)), there exist outcomes A, A+ and B such that it is indeterminate whether $V(A)$ is greater than $V(B)$, it is indeterminate whether $V(B)$ is greater than $V(A)$ and it is indeterminate whether $V(A) = V(B)$, where "it being indeterminate that p" implies that it's not the case that it's true that p. I will be focusing on the first version, but everything I say would apply equally well to the second.

⁴⁷ Once again, if you are skeptical of assigning values to outcomes as a whole, you can think of these claims as being about the value to some person (or group of people). Then MUSHY WORLD says that our insensitivity to sweetening in cases in which we are choosing between two ways of helping a person results from the fact that the values *for that person* of the two outcomes are on a par.

and of itself, is not the kind of thing that can be measured precisely. According to MUSHY HEAD, the incommensurability we experience is not a result of the evaluative nature of the world; rather, the world comes with a standard value function which assigns to every outcome a unique real number. Our insensitivity to sweetening is due to an epistemic limitation on our part; we simply do not know what this function is.⁴⁸ An agent who was both perfectly rational and whose preferences were completely aligned with the evaluative nature of the universe would never be insensitive to sweetening.

I will now argue that if we accept MUSHY HEAD, we can solve the problem. Recall that according to the MUSHY HEAD view, a fully rational agent who knew all there was to know about the future, and knew all the facts about value, would know which charity to donate to. I am going to call agents that are perfectly rational *and* know all the facts there are to know about value "Perfect Agents". The solution to the problem relies on the fact that, if we accept MUSHY HEAD, we need to distinguish between two measures of value. One is the measure given by the *correct* value function, the one that represents the true evaluative nature of the world. The other is given by the set of value functions that *you* use as a result of your ignorance of the correct value function. Since we have two ways of measuring value, we also have two ways of measuring expected value. For any action, we can consider the expected value of that action according to the correct value function and the expected value of that action according to your set of value functions.

The basic idea is this: CHOICE REFLECTION tells you to defer to future expected value judgments. But, if we believe MUSHY HEAD, the following question will arise: what *kind* of future expected value judgments should you defer to? Should you defer to future expected value judgments that are based on the *correct* expected value function? Or, should you defer to future expected value judgments that are based on the expected values you will assign as a result of your ignorance of the correct value function? As I will argue, while you should always defer to future expected value judgments based on the *correct* value function, you should *not* always defer to the future expected value judgments that you will make. If MUSHY HEAD is right, although you know that, in the future, you will be indifferent between going east and going west,

⁴⁸ Note that the kind of ignorance we have about which value function is correct is not the kind that can be represented by precise probabilities (if we were ignorant about which function was correct, but could assign a precise probability to each possible function being the correct one, we would still never be insensitive to sweetening). I discuss imprecise probabilities in Schoenfield (forthcoming).

you *don't* know that this future indifference of yours will be the attitude warranted by the correct value function. So the right version of CHOICE REFLECTION, as I will argue, will not tell you to defer to this future judgment. This is why there is no reflection worry for the MUSHY HEAD theorist.

Here is an outline of the argument:

- (1) In CHARITY OPAQUE, you do *not* know that your future assignments of expected value will be based on the *correct* value function.
- (2) The right version of the reflection principle, according to a MUSHY HEAD theorist, will only tell you to defer to future *correct* expected value judgments.
- (3) The right version of the reflection principle, according to a MUSHY HEAD theorist, will not tell you defer to your future judgment in CHARITY OPAQUE. (1,2).

Defense of (1):

I will begin by defending the first premise: the claim that, in CHARITY OPAQUE, although you know what expected value judgments you will have in the future, you do not know that these expected value judgments are based on the correct value function. Let's use " $EV_{correct}(X)$ " to refer to the expected value of X according to the correct value function and " $EV_{you}(X)$ " to refer to the expected value of X according to the set of value functions that represent your evaluative attitudes. In CHARITY OPAQUE, before you know which charity is in which location, you know that $EV_{correct}(East)$ is greater than $EV_{correct}(West)$. Why? Your set of value functions includes a range of functions, including (hopefully!) the correct value function. As argued earlier, the expected value of going east is greater than the expected value of going west according to *every* value function in the set. So, in particular, the expected value of going east is greater than the expected value of going west according to the correct value function. However, while you now know that $EV_{correct}(East)$ is greater than $EV_{correct}(West)$ and you know that, later, you will be indifferent between East and West, you do not know that, later, you will judge $EV_{correct}(East)$ to be no greater or less than $EV_{correct}(West)$. After all, once you know which charity is in which location, the value functions in the set will disagree about which location has greater value. Since you do not know which function is the correct one, you will not know what the relative expected value of going east and west is according to the correct value function.

Defense of (2)

I will now argue that, if you accept MUSHY HEAD, the correct principle will only tell you to always defer to future *correct* expected value judgments. Here is the right version of the principle for a MUSHY HEAD theorist:

CHOICE REFLECTION MUSHY HEAD: Suppose you are faced with a choice between doing X and doing Y. If you know that at a later time you will rationally, without loss of information, judge that $EV_{\text{correct}}(X)$ bears a certain relation R (greater than, less than, no greater and no less than) to $EV_{\text{correct}}(Y)$, you should now judge that $EV_{\text{correct}}(X)$ bears R to $EV_{\text{correct}}(Y)$.

Why should we only accept the version of the principle which tells you to always defer to future *correct* expected value judgments? Because, if you accept a reflection principle that tells you to always defer to *your* future expected value judgments, even when they are not based on the correct value function, the principle would yield the result that sometimes you should make expected value judgments that differ from *those you know to be correct*. This would be a bad result, since in cases in which you are concerned with maximizing expected value, what you should be concerned with is maximizing expected value based on the value function that correctly represents the facts about value in the world! According to MUSHY HEAD, we are not always able to make correct expected value judgments, but when we are, these are surely the ones that should guide our actions.

To see why a reflection principle that would always have you defer to your future expected value judgments would yield the absurdity mentioned above, let's consider an agent who is perfectly rational and knows what the correct value function is. Call this agent the "Perfect Agent." How would the Perfect Agent respond to CHARITY OPAQUE? Since we know that, in the opaque case, $EV_{\text{correct}}(\text{East})$ is greater than $EV_{\text{correct}}(\text{West})$, we know that the Perfect Agent would, at least initially, judge that she ought to go east. Crucially, this would be the case even if she knew about *your* future indifference between going east and going west. Here is why: The Perfect Agent will not always defer to your future expected value judgments because your judgments are not always based on the correct value function. Nonetheless, there might be times in which she will take your future judgments into account when forming her own because your judgments may be based on more information than hers. The reason she might want to take your future judgments into account when they are based on more information is that she might be able to use your judgments to infer facts about what her judgments would be were she to have the same information that you will have. In the case we are discussing, however, the Perfect Agent *learns nothing* about what her future judgments would be, by learning about your future judgments. This is because her future attitude will depend on which charity is in which location,

but you will be indifferent *regardless of what you find out*, so learning that you will be indifferent gives her no information about which charity is where and, hence, no information about what her future attitude will be.

The following analogy may be helpful: Suppose I prefer red marbles to green marbles. There is a marble in the right box and a marble in the left box but I don't know which marble is in which box. You, however, can see which marble is in which box and form a preference. If I know something about your preferences, even if they are different from my own, learning which marble you prefer can give me information about which marble I would prefer. However, if I know that you are indifferent between red and green marbles, learning about your preferential attitudes once you see which marble is in which box is completely useless to me. And it would be silly of me to be guided by your preferences, even though they are more informed than my own. Similarly, since you will be indifferent no matter what happens, the Perfect Agent has no reason to defer to your future judgments in this case.

Since you know that the Perfect Agent with the same information you have would be unmoved by your future indifference, you know that, in the opaque case, $EV_{\text{correct}}(\text{East})$ is greater than $EV_{\text{correct}}(\text{West})$. Since you are trying to maximize expected value according to the correct value function, whenever you know what choice the correct value function dictates, this is the choice that you should make. This is true even if you know that you will make a different judgment in the future. Thus, the right version of the reflection principle, according to a MUSHY HEAD theorist, will not tell you to defer to your future judgment in CHARITY OPAQUE.

In sum, if MUSHY HEAD is right, there is no reflection problem. We get the result that you should help the weary traveler and this result isn't in conflict with any important principles of decision making. However, MUSHY HEAD faces its own difficulties. For, according to MUSHY HEAD, the world comes with *incredibly* precise facts about value. This seems very implausible, and so I think that this solution should be rejected. However, it is worth noting that MUSHY HEAD is the view that traditional utilitarians have embraced, and so those that are comfortable with the existence of a "magic value function in the sky"⁴⁹ that assigns to each outcome a unique real number representing its value, may find this solution to be satisfactory.

⁴⁹ I first heard the term "magic X function in the sky" at a talk given by Brian Weatherson at MIT in the Spring of 2011.

Since I am not, I will now propose a solution that avoids the metaphysical commitments that MUSHY HEAD faces.

5. The Acquaintance Based Solution

We have seen that, if we maintain MINIMAL CONSEQUENTIALISM, we are faced with a problem. If we think that incommensurability is a result of the evaluative nature of the world, it is difficult to see why we should help the traveler. If we think that incommensurability is a result of ignorance on our part, we don't face the same problem, but this view is extremely implausible. I think that the way out is to deny MINIMAL CONSEQUENTIALISM.

It was stipulated in the charity cases that, no matter what choice you make, you will not be violating someone's rights, failing to meet some associative obligation, lying, failing to respect your own interests and projects, or any of the standard no-nos that many non-consequentialists frequently appeal to in explaining why consequences aren't all that matters. Nonetheless, it seems that being moral requires making one choice rather than another despite the fact that the value of one outcome is no better than that of the other. In this section, I argue that, if we take facts about our *acquaintance* with the reasons to help others to be relevant to the question of how we should help, we can explain why this is so.

Here is the plan: First, I will say a little bit about what I mean by acquaintance. Second, I will show that, in the cases we have been discussing, acquaintance considerations provide a reason to help the traveler. I will also show why it is that, once we take acquaintance seriously, the idea that we should defer to our future judgments is misguided. Finally, I will illustrate that taking acquaintance seriously may help explain some other puzzling issues in ethics.

Let me begin by saying what I mean by acquaintance with the reasons to help others. The idea is as follows: There are more or less intimate ways of knowing the reasons to help someone. In the beginning, I suggested as an example the case of Carrie who needs 100 dollars. Knowing that Carrie needs 100 dollars gives us reason to help Carrie. We might also know that Carrie needs 100 dollars for medicine. Or maybe we even know that Carrie needs 100 dollars for medicine for her asthma. It's not as if Carrie needing money, Carrie needing money for medicine, and Carrie needing money for asthma medicine are three independent reasons to help Carrie. However, learning the additional information about Carrie, makes us better *acquainted* with the reasons to help her.

I am going to suggest that, all else equal, the greater our acquaintance with the reasons to help others, the greater our obligation to help them. Why should we think this is the case? Here is the rough idea: Being a good person requires being sensitive to the needs of others and being moved by these needs to help. The greater our acquaintance with the reasons to help others, the more *insensitivity* is displayed by ignoring them. Since insensitivity is a moral failing, failing to help in situations in which you are very well acquainted with the reasons to help is worse than failing to help in situations in which you are not as well acquainted.

Before providing some further motivations for taking acquaintance seriously, let's see why it is that, if acquaintance does matter, we can solve our problem. In CHARITY OPAQUE, you lack knowledge of which charity is in which location and, as a result, you are not well *acquainted* with the reasons to go east and the reasons to go west.⁵⁰ You are, however, well acquainted with one reason to go east, and this is the fact that by doing so you could help the traveler. If you are a sensitive person, having been acquainted with the traveler's suffering, you will simply be moved to help him.⁵¹ Failing to help the traveler is insensitive, but not because the outcome in which you help the traveler is better than the outcome in which you do not. You know that neither outcome is better than the other. However, your acquaintance with the traveler's needs gives you a reason to go east, rather than west, that is independent of the values of the outcomes.

⁵⁰ There is a sense in which you are acquainted with the reasons to go east and west even in the opaque case. For you know that by going east you will help the meal program *or* the arts program and you know that by going west you will help the meal program *or* the arts program. As we'll see, however, taking acquaintance seriously requires giving more weight to the reasons that are *non-disjunctive*. If we allowed acquaintance with disjunctive reasons to count, none of the cases in which it seems intuitive that something like acquaintance should matter would actually be cases in which there was a difference in the degree of acquaintance between the two options. For whenever you're in ignorance about the reasons to do some action, you are always acquainted with some enormous disjunction which includes all the possible reasons to do the action.

⁵¹ There is another way in which acquaintance might be relevant. You might think that the greater acquainted you are with someone's needs the more *disrespectful* it is to ignore those needs. If all you are acquainted with is the needs of the traveler, then failing to meet his needs is disrespectful to him. Since you don't know which charity is in which location, failing to help him suggests that you are not taking his needs seriously. If, however, you are acquainted with the needs of the charity in the east and the charity in the west, failing to help the traveler is not disrespectful because your failure to help needn't reflect a failure to take his needs seriously. It just means that there were graver needs that you were taking seriously. The reason I don't find this account satisfying is the same as the reason I don't find the "disrespect" explanation for insensitivity to sweetening plausible in the transparent case; namely, that such an explanation only works in an *interpersonal* case. However, the same issues arise in *intrapersonal* cases. For example, if we knew that by going east we could help the school one way and by going west we could help the school another way, and we also allowed the sweetening to be a benefit to the school, it would not make sense to say that by taking one option rather than another we were showing disrespect to the students or not taking their needs seriously, since whatever we will be doing, will involve helping the school.

If we acknowledge the importance of acquaintance, a principle which tells us to always defer to the judgments of our later selves is unmotivated. For although our later selves are the experts when it comes to maximizing the value of the *outcome*, our later selves are acquainted with different facts than our current selves. Since our degree of acquaintance is part of what determines the degree of our obligation to help, we should not always defer to our later judgments. Knowledge of future acquaintance is not the same as acquaintance and so, it is the facts about our *current* acquaintance that determine what we should do. In the opaque charity case, you know that, in the future, you will be acquainted with a host of reasons to go west that you are not acquainted with now, but what's important is that you are not acquainted with them *now*. In the transparent case, in contrast, you know a lot about the charities, and so you are acquainted with the reasons to help each of them. This is why, in the transparent case, you are not obligated to take the sweetened option. For in the transparent case, not only do you know that there are reasons to take the non-sweetened option, but you are *acquainted* with those reasons.

I have argued that, if we take acquaintance seriously, we can solve the problem, but you might think that acquaintance considerations are not relevant to what one *ought* to do. Perhaps, instead, you think that all that the considerations above suggest is that, if you fail to help the traveler, you have revealed a bad character trait. Nonetheless, one might say, this doesn't imply that you *ought* to help the traveler. This is certainly a viable option, and I have no knock-down argument against it. However, it is not my preferred way of treating the case. I think that the considerations mentioned earlier suggest that you really *should* help the traveler, and not only that doing so reveals a bad character trait. However, I also think that the explanation of the wrongness of the act is closely connected to the problematic character trait that would be revealed in ignoring the traveler.

In the remainder of the section, I will describe some examples which suggest that taking sensitivity, and hence, acquaintance, seriously can explain some other judgments in the ethics of helping. Here is the first example: I spent some time volunteering at an organization in Thailand which rehabilitates Thai elephants that have been abused or neglected. I think that, after learning about all the horrible things that happen to these elephants, I had an obligation to support such an organization in one way or another (volunteering, donating money, etc). In fact, I think that given the choice to support the Thai elephant organization or some other wonderful organization that helps, say, the tigers in Kenya, I ought to help the elephant cause. In contrast, I do not think

you have an obligation to help the Thai elephants rather than the Kenyan tigers. The reason I should help the elephants has nothing to do with the fact that the outcome brought about by helping the elephants would be better than the outcome brought about by helping the tigers. For I know that if I knew more about the tiger situation in Kenya, I would be equally moved to help them. In fact, I may even know that, in the future, I *will* learn about the Kenyan tigers and, at that time, I will judge that helping the elephants was no better than helping the tigers. The reason I should help the elephants *now* is that being moral requires being sensitive to certain kinds of considerations. Failing to help the elephants when I know so much about their suffering would be more insensitive than failing to help the tigers, whose situation I know less about.

It is not just detailed knowledge of someone's needs that I think can make a moral difference. There are different ways of being acquainted: knowledge is one, but seeing, for instance, may be another. The second example I will consider is Peter Singer's famous case of the child drowning in the pond. Singer argues as follows: if we see a child drowning in a muddy pond, we ought to go save the child even if it means ruining our expensive shoes. Therefore, for exactly the same reasons, we ought to donate money which would otherwise be used for morally insignificant things (like expensive shoes), to save the lives of people abroad. Singer takes it for granted that there is no morally significant difference between the pond case, and the case of the distant needy. Here is what he says:

"I do not think I need to say much in defense of the refusal to take proximity and distance into account. The fact that a person is physically near to us, so that we have personal contact with him, may make it more likely that we shall assist him, but this does not show that we ought to help him rather than another who happens to be further away. If we accept any principle of impartiality, universalizability, equality, or whatever, we cannot discriminate against someone merely because he is far away from us (or we are far away from him)" (232).

I by no means intend to argue that we have no obligations to the distant needy. However, I do think that there is a problem with Singer's argument about the relevance of proximity. Singer says that if we accept any principle of impartiality, universalizability, equality or whatever, we cannot discriminate against someone because he is far away from us. Here is the claim in the neighborhood that I agree with: If we accept any principle of impartiality, universalizability, equality or whatever, we should not judge the life of someone nearby to be *more valuable* than that of someone far away. Or, similarly, we should not think of the *outcome* of saving the near as

any better than the *outcome* of saving the distant. However, I think you *should* reason to save the drowning child, and the reason for this is that failing to do so would be more insensitive than failing to save a child far away. It is worth noting, however, that my explanation of the case doesn't make mere proximity morally relevant. Proximity is only relevant in so far as it increases acquaintance. We can easily imagine a case where you are more acquainted with the reasons to help a distant child than a nearby child, and in such a case, all else equal, you should help the distant child.

Singer grants that we may be more psychologically disposed to save the drowning child than the distant needy but thinks that this has no bearing on what we ought to do. I think that this is too quick. The reason we are more disposed to save the drowning child is that the suffering of the child is vivid to us. The reasons to sacrifice our shoes, in this case, are, so to speak, staring at us in the face. If you are a decent person, you will be moved by the suffering of the child and will sacrifice your shoes without thinking twice. Failing to empathize with and be moved by the suffering of a child drowning in front of your eyes is a serious defect. Failing to empathize with and be moved by suffering that you are not as well acquainted with may be a defect, but it is certainly not as serious.

Consider an analogy: Being a good cook requires being sensitive to subtle differences in flavor. Perhaps perfect cooks are maximally sensitive so that even extremely subtle differences in flavor will be detected. However, if the differences in flavor are significant enough, even a minimally decent cook should be able to detect them. Acquaintance is analogous to the strength of flavor. When we are well acquainted with reasons to help, and we fail to help, we may fail to be even minimally decent. The better we are, the more sensitive we will be to reasons that we are not very well acquainted with. However, just as development of taste may be a process, development of this kind of moral sensitivity may be a process as well. Perhaps, as we become more sensitive, the importance of acquaintance diminishes. Perhaps the ideally moral agent will be very moved by the suffering of a distant child that she knows absolutely nothing about. I do think, however, that even if the importance of acquaintance diminishes, it does not entirely disappear.

I would like to end with one other way in which the moral relevance of acquaintance can be helpful. Susan Wolf (1982) argues that there is something unattractive about the kind of person who is, what she calls, "a moral saint." Moral saints are people "whose every action is as

morally good as possible” (419). The moral saint, Wolf argues, will tend to be dull, humorless, and lack appreciation for some of life’s great delights. Although there is certainly something unattractive about being dull and humorless, some of the features that Wolf ascribes to moral saints are features that presuppose that acquaintance is not morally relevant. For example, Wolf describes her moral saint as someone who devotes all of his time to “feeding the hungry, or healing the sick, or raising money for Oxfam” and one of the many things that the moral saint misses out on, according to Wolf, is an appreciation of literature. Indeed, if you spend all of your time raising money for Oxfam, you probably will not have time to read *Pride and Prejudice*. However, if acquaintance is morally relevant, the moral saint will not spend all of his time raising money for Oxfam. This is because he may have greater acquaintance with the needs of a child in the neighborhood who needs help in English literature. Again, I am not saying we don’t have obligations to the distant needy, and I am certainly not suggesting that acquaintance is the *only* consideration relevant to questions of whom to help. However, I think that taking acquaintance considerations to have *some* moral relevance is sufficient for pointing to a problem with Wolf’s argument.

The following quote from Wolf describes the view of morality that makes sainthood look unattractive:

“The moral point of view, one might say, is the point of view one takes up insofar as one takes the recognition of the fact that one is just one person among others equally real and deserving of the good things in life as a fact with practical consequences, a fact the recognition of which demands expression in one’s actions and in the form of one’s practical deliberations.” (436-437)

Indeed, if we think it follows from “everyone is equally deserving of the good things in life” to “we must treat everyone’s needs equally,” a moral saint will be the kind of person Wolf describes as unattractive. If, however, we take acquaintance seriously, the recognition of the equal importance of humans needs need not translate into an obligation to treat those needs equally.

6. Conclusion

Plausibly, the facts about value are not precise. That is, the facts about value cannot be represented by a standard value function which assigns to each outcome a unique real number. I have argued that the imprecision that we find in value poses a problem. The problem arose by

considering cases in which, seemingly, our only concern is the value of the outcome. If we are concerned with maximizing the expected value of an outcome, we ought to defer to the value judgments of our future, more knowledgeable selves, since they are in a better position than we are to bring about the desired outcomes. However, I suggested that, even when standard non-consequentialist constraints are satisfied, the outcome is not all that is at stake when we are determining whom to help. Moral people are people who are marked by a sensitivity to certain kinds of considerations and who are disposed to be moved by the suffering of others. Since part of being moral requires being sensitive, it should not be surprising that our obligation to help others depends, in part, on our degree of acquaintance with the reasons to help them.

References

- Broome, J. (1997) "Is Incommensurability Vagueness?" In R.Chang (ed) *Incommensurability, Incomparability, and Practical reasoning*. Harvard University Press: Cambridge, MA. pp. 67-89
- Chang, R. (1997) Introduction to *Incommensurability, Incomparability and Practical Reason* Harvard University Press: Cambridge, MA.
- Constantinescu, C. (forthcoming) "Value Incomparability and Indeterminacy." *Ethical Theory And Moral Practice*.
- Christensen, D. (2004) *Putting Logic in Its Place: Formal Constraints on Rational Belief*. Oxford, UK: Oxford University Press.
- Doggett, T. (2009). "What is Wrong with Kamm and Scanlon's Argument Against Taurek." *Journal of Ethics and Social Philosophy* 3(3).
- Graham, P. (2010) "In Defense of Objectivism about Moral Obligation." *Ethics* 121 (1).
- Hare, C. (2010) "Take the Sugar." *Analysis* 70 (2): 237-247.
- Kamm, F. (2000). "Nonconsequentialism" in Lafollette, Hugh, ed., *The Blackwell Guide to Ethical Theory*. Malden, MA: Blackwell.
- Raz, J. (1985) "Value Incommensurability: Some Preliminaries." *Proc. of the Aristot. Soc.* 86: 117-134
- Schoenfield, M. (forthcoming) "Chilling Out on Epistemic Rationality: A Defense of Mushy Credences." *Philosophical Studies*.

Singer, P. (1972) "Famine, Affluence and Morality." *Philosophy and Public Affairs* 1 (3): 229-243.

Wolf, S. (1982) "Moral Saints." *Journal of Philosophy* 79 (8): 419-439.

Chapter 3

Permission to Believe

Why Permissivism is True and What it Tells Us about Irrelevant Influences on Belief

1. Introduction

I was once talking to a very religious friend (let's call her "Carmen") about whether or not her particular religious beliefs were justified. As these conversations tend to go, we each proposed arguments that challenged the other's beliefs', responded to them, deemed the other's responses unsatisfactory, and neither of us budged. As a last attempt, I said to her: "Look, you must realize that if you had grown up somewhere else, you would not have all of these beliefs. You only believe as you do because of the influence of the people around you. How, then, can you be so sure you are right?" This thought, which had clearly occurred to her in the past, she found deeply troubling. I have another friend (let's call him "Joe") who once told me that the only challenge to his faith that ever *really* concerned him was the fact that his beliefs were caused by the community he was raised in.

I find the concern that Carmen and Joe share intriguing and I think it is one that many of us can relate to. It can seem very worrying that many of our deeply held convictions were caused by seemingly irrelevant influences, like the community we grew up in, the school we went to, or the friends we hang out with. My primary aim in this paper is to offer a response on my friends' behalf; that is, to argue that in many of the cases in which people worry about the irrelevant influences on their beliefs, they need not.

Given how pervasive these kinds of influences on belief are, I think it is extremely important that we figure out how we should respond when we learn about them. For if it turns out that, upon learning that a belief was caused by an irrelevant influence, we are rationally required to give up this belief, the consequences would be quite drastic. This is because many of the beliefs that have been caused by irrelevant influences are religious, moral, political and philosophical beliefs. These are the kinds of beliefs that are very central to who we are, and to important decisions that we make about how to structure our lives.

Here is the general plan for the paper: first, I will provide some motivations for *permissivism*: the claim that sometimes, there is more than one rational response to a given body of evidence. In recent years, permissivism has fallen out of favor and so it will take some work

to defend this view and respond to the compelling arguments that have been offered against it. Second, I will argue that, if we accept permissivism, the concern raised by irrelevant influences on belief is, in many cases, unwarranted. Third, I will raise a problem for my view and respond to it. And finally, as a bonus, I will show how the view I defend helps sort out some issues related to peer disagreement.

Before delving into the argument, I would like to provide a more nuanced characterization of the view I will be defending, but first, I need to introduce some terminology. I'm going to call cases where we learn that our beliefs were caused by an irrelevant factor "IF cases." Later in the paper it will become clearer exactly what kind of factor counts as "irrelevant." So, for now, I hope you have an intuitive sense for which sorts of causes of beliefs are the worrisome ones. Also, I am going to use the term "doxastic attitude" to refer to any attitude which measures the degree of confidence an agent has in a proposition. Believing that p , being agnostic about whether p , and having credence 0.78 in p , are all examples of doxastic attitudes.

In a nutshell, my view is this: there are non-permissive cases (cases in which there is only one rational way to respond to a body of evidence) and permissive cases (cases in which there is more than one rational way to respond to a body of evidence). If we start out in a non-permissive case, and then we learn that we are in an IF case, a significant reduction of confidence may indeed be warranted. However, this is not so in permissive cases. I am going to argue that if you believe p in a case in which, given your evidence, it is permissible to believe p but also permissible to believe $\sim p$, then gaining additional evidence which suggests that your belief that p was caused by an irrelevant factor should not worry you. (Put in terms of defeaters, the rationality of your belief in p is not defeated by the new evidence you get about the cause of your belief). You might find this obvious. After all, if I am a permissivist, and think that given my original body of evidence E , both belief and disbelief in p are rational, why should learning about the *cause* of my belief matter at all? I ended up with a rational belief! Isn't that all that matters? Although tempting, I will argue that this line of thought should be rejected. For this line of thought requires the permissivist to think of her beliefs in permissive cases as, in a sense, arbitrary, and, as I will show, a plausible version of permissivism will not yield this result. The argument for the claim that in permissive cases we should not be concerned with irrelevant influences will, therefore, require some care. In fact, it will turn out that there are very intricate

connections between the debate about permissivism and the debate about irrelevant factors. Some of the arguments against permissivism appeal to certain judgments about IF cases, and some of the motivations for reducing confidence in IF cases rely on views about permissivism. Part of what I aim to do is untangle these knots and make clear exactly what the connection between these two debates is.

Here is how the argument will go:

Main Argument

P1. Permissivism is true.

P2. If permissivism is true, the view that you should reduce confidence significantly in permissive irrelevant influence cases is unmotivated.

C1. You don't need to reduce confidence significantly on the basis of irrelevant influences in permissive cases. (P1, P2).

In section two I defend the first premise, and in section three I defend the second premise.

2. Permissivism (Defense of P1)

In this section I motivate the first premise of the main argument.

2.1 Motivations for Permissivism

Consider the following claim:

UNIQUENESS: For any body of evidence E, and proposition P, there is only one doxastic attitude to take towards P that is consistent with being rational and having evidence E.

Permissivism is the denial of UNIQUENESS. It is the claim that there are some cases in which there is more than one rational response to a given body of evidence.⁵² Before describing the motivations for permissivism, I should note that the debate about permissivism is different from (though related to) the debate about how to respond to disagreement. There is one question about whether two people who evaluate the same evidence can rationally come to different

⁵² Actually, the view that I will be defending here (and calling “permissivism”) is somewhat stronger than the denial of UNIQUENESS. For I will be maintaining that there are some cases in which it is rational to believe p given E and also rational to believe ~p given E. However, it is consistent with denying UNIQUENESS that no such cases exist. You might adopt a moderate version of permissivism according to which there are sometimes multiple permissible attitudes to take towards some proposition p, but it is never the case that given some body of evidence E, it is rational to believe p on the basis of E, and rational to believe ~p on the basis of E. However, for reasons I cannot get into here, I do not find this moderate position to be a well-motivated one. (The basic idea is that the moderate position will face all the problems that more extreme versions of permissivism face, without getting the benefits that the more extreme versions enjoy). I will therefore be setting this view aside for the remainder of the paper.

conclusions (this is what the permissivism debate is concerned with) and a different question about what two people should do when they *learn* that the other has come to a different conclusion (this is what the disagreement debate is concerned with). This section is concerned only with the first question.

Permissivism is motivated by both intuitive and theoretical considerations.

(1) Intuitive Motivations

Gideon Rosen (2001) gives the following example in his defense of permissivism:

“It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence...Paleontologists disagree about what killed the dinosaurs. And while it is possible that most of the parties to this dispute are irrational, this need not be the case. To the contrary, it would appear to be a fact of epistemic life that a careful review of the evidence does not guarantee consensus even among thoughtful and otherwise rational investigators” (71).

It is not just in scientific contexts in which it seems that people can reasonably arrive at different conclusions on the basis of the same body of evidence. We can imagine cases in which members of a jury are examining a complex body of evidence about who committed a crime, or people considering the evidence for and against the effectiveness of acupuncture, or the existence of God. In these cases too it seems that people may rationally come to different conclusions on the basis of a shared body of evidence.

There is a way of thinking about evidence which suggests that the cases that Rosen and I are appealing to do not, in fact, motivate permissivism. Alvin Goldman (2010), for example, points out that two people rarely share *all* of their evidence. Consider, for example, the jury case. Although all the members of the jury know the same facts about this particular crime, they have all been exposed to different evidence throughout their life. Goldman argues that these differences may be relevant to the question of what it is rational to believe about the suspect. If Goldman is right, we might be able to maintain UNIQUENESS and acknowledge that the members of the jury can reasonably come to different conclusions. This is because we might be able to explain the diverse conclusions by appealing to the fact that the jury members each had a different body of evidence to evaluate. But to motivate permissivism, we need cases where people with the *same* evidence reasonably come to different conclusions. Therefore, you might think, these cases do not, in fact, motivate permissivism.

There are two things worth noting here: first, it seems right that, in *some* cases, diversity of opinions can be accounted for by a diversity of evidence. But even if this approach works for some cases, it does not seem plausible that *all* cases in which, seemingly, people rationally come to have differing opinions on some matter can be explained by a diversity of evidence. Take, for example, Rosen's case of the disagreement among paleontologists about what killed the dinosaurs. It seems unlikely that this disagreement can be rationalized by appealing to subtle differences in the bodies of evidence that the paleontologists have. Although the paleontologists may have grown up in different places, had different childhood experiences, and shop at different grocery stores, presumably, the evidence relevant to the question of who killed the dinosaurs is limited primarily to the facts paleontologists study in their academic lives. This evidence, which they get from academic books and papers, is the kind of evidence that can be easily shared. Thus, it is implausible that we can rationalize all of the paleontologists' disagreements by appealing to differences in their total body of evidence.

The second point worth noting is that using a suggestion like Goldman's to explain away all the cases we have been considering (which I do not think was actually Goldman's intention), comes at a significant cost. Here is why: a crucial part of our epistemic life involves deferring to the opinions of others who are better placed to answer certain questions than we are. Normally, we only know rather *rough* facts about other agents' bodies of evidence. However, the differences that Goldman would need to appeal to explain away all of the cases of seemingly rational disagreement are extremely *fine*. If what it is rational to believe is extremely sensitive to incredibly fine differences in evidence, it will be very difficult to use our merely rough knowledge of others' evidential states to make judgments about whom to defer to.

In sum, there are two reasons to think that Goldman's appeal to subtle differences in evidence cannot undermine the intuitive motivation for permissivism. One reason is that it seems implausible to appeal to subtle differences in evidence to explain *all* cases of seemingly rational disagreement. The second reason is that attempting to make too much of these subtle differences comes at a significant theoretical cost: it makes it difficult to explain how we can frequently do a good job at figuring out who is best placed to answer particular questions.

(2) Theoretical Motivations

Certain plausible theories of justification require the truth of permissivism (in fact, according to Igor Douven (2009), *all* current plausible theories of justification require

permissivism!).⁵³ Coherentism,⁵⁴ conservatism,⁵⁵ and subjective Bayesian,⁵⁶ are perhaps the most well known permissive theories. I will not be able to go into the motivations for these theories here, but I do want to note that at least some theories of justification that reject permissivism are laden with some unfortunate metaphysical commitments. Let us consider, as an example, someone who thinks in terms of degrees of belief (like a Bayesian), and thinks that the appropriate way to respond to evidence is with a *credence* (a real number between 0 and 1 that measures one's confidence in a proposition). On such a picture, given any body of evidence E, and proposition p, there is unique real number which represents the rational response to E. Consider the proposition that there exist more than three hundred elephants. According to the picture we are considering there is a unique real number that is the appropriate credence for me to have in this proposition.⁵⁷ It is, however, somewhat mysterious what could ground such facts. It is not as if there exist some general principles like "infer to the best explanation" or "believe what is entailed by your evidence" or "set your credences to the known objective chances" that can determine a unique real number measuring the reasonable credence to have in such a proposition. (Carnap (1950) heroically attempted to come up with a set of principles that would do this sort of thing, but the Carnapian project ultimately failed).⁵⁸ The defender of UNIQUENESS, then, will have to take such facts to be brute, and the commitment to a preponderance of these brute facts seems to be an unattractive feature of the theory.

Gideon Rosen seems to be motivated by a similar concern when discussing a case which he takes to be one of reasonable disagreement between us and a group of people called "The Bedrockers."

Are the Bedrockers unreasonable? If they are, then it should be possible to locate their mistake: to describe a compelling dialectical route... or at the very least to identify some principle or rule of inference which they reject, the rejection of which strikes us—when we hold it up to the light—as somehow crazy or silly or perverse or unintelligible.

⁵³ For a thorough discussion of which views are ruled out by UNIQUENESS see also Ballantyne and Coffman (forthcoming).

⁵⁴ See, for example, Quine and Ulian (1970), Bonjour (1985) and Elgin (1996).

⁵⁵ See, for example, Harman (1986) and Lycan (1988).

⁵⁶ See, for example, Jeffery (1965).

⁵⁷ Some people, like Levi (1974) and Joyce (2005) think we can avoid such problems by moving to a version of Bayesianism that represents doxastic states with a *set* of probability functions, rather than a single function. However, I do not think that moving from single functions to sets of functions actually solves the problem. (This version requires that that for any body of evidence there is a unique *interval* which represents the credence one ought to have in a proposition). Additionally, there are independent reasons to be worried about this version of Bayesianism. See White (2010)^a, Elga (2010), and Schoenfield (forthcoming).

⁵⁸ See, for example, Burks (1953), Horwich (1982), and Titelbaum (2010).

Rosen is suggesting that, in many cases, if one of the two positions is, in fact, unreasonable, this cannot be a matter of brute fact. Some *principle* is necessary to explain what is unreasonable about one of these positions. But plausibly, as the cases we have been discussing suggest, general principles of reasoning will not ground facts about justification that will determine a unique reasonable attitude to take in any given case.

Here is another way of making the point: rationality can be thought of as a kind of capacity. Somebody committed to UNIQUENESS thinks that the possession of this capacity will allow one to make extremely fine grained distinctions. Consider, for example, the proposition that all ravens are black. The UNIQUENESS defender is committed to thinking that there is some number n , such that if you see n or more ravens, all of which are black, it is rational to believe⁵⁹ that all ravens are black, but if you see fewer than n , it is not rational to believe that all ravens are black.⁶⁰ It is, however, implausible that possessing rational capacities would allow one to recognize a unique such number! Rational capacities are capacities that one can apply to a broad range of epistemic situations. They allow one to make computations and inferences, to recognize certain arguments as good or bad, to recognize certain features of evidence as telling in favor or against theories, and so forth. Being in possession of such capacities may enable us to *rule out* certain patterns of attitudes in response to observing black ravens, but it hard to see how such capacities would allow one to determine that you need a minimum of 28 ravens to warrant believing that all ravens all black.

In sum, at least on some theories of justification, the commitment to UNIQUENESS has significant disadvantages. This not to say that there is no theory of rationality that sits well with the UNIQUENESS assumption. My point is only that there are some *prima facie* theoretical reasons to reject UNIQUENESS. These reasons, combined with the intuitive motivations I described above are, I think, enough to support the claim that, barring any serious reasons to reject permissivism, we should embrace it. In the next section, I will respond to some of the arguments that have been put forward against permissivism.

⁵⁹ Or, if you prefer, have a credence above 0.9, or have a credence interval above 0.9.

⁶⁰ Of course, the rationality of believing that all ravens are blacks depends on more than the number of ravens you see, but let's hold everything else fixed for now and consider only differences in the number of observed black ravens.

2.2 Problems with Permissivism

Despite the intuitive and theoretical motivations for permissivism, many have thought that permissivism is not a viable option. Some of the recent defenders of UNIQUENESS include Roger White (2005), David Christensen (2007), Richard Feldman (2007) and Ernest Sosa (2010). I will respond to their arguments below. But before doing so, I should note that what I say here will not constitute a defense of *all* versions of permissivism. I think the arguments for UNIQUENESS point to a real and serious problem with certain versions of permissivism. However, what I take to be the best version of permissivism is immune from these arguments. According to the version of permissivism I will be defending here, what one ought to believe depends, in part, on what epistemic standards one has. On this view, if two people with the same evidence reasonably have different opinions about whether *p*, it is because these people have each adopted a different set of reasonable epistemic standards.

What are an agent's epistemic standards? There are different ways of thinking of epistemic standards. Some people think of them as rules of the form "Given *E*, believe *p*!"⁶¹ Others think of them as beliefs about the correct way to form other beliefs.⁶² If you are a Bayesian, you can think of an agent's standards as her prior and conditional probability functions. Since what I will be saying does not rely on a particular understanding of what a standard it is, we can just think of a set of standards as a function from bodies of evidence to doxastic states which the agent takes to be *truth conducive*. Roughly, this means that the agent has high confidence that forming opinions using her standards will result in her having high confidence in truths and low confidence in falsehoods.⁶³ On the version of permissivism that I will be defending, there are multiple permissible epistemic standards, and what makes it permissible for agents to have different doxastic attitudes is that different attitudes may be prescribed by their different standards. In responding to the arguments against permissivism, this will turn out to be important.

I will first present an argument against permissivism which I think simply does not speak to the version of permissivism I am defending. I will then discuss a cluster of worries that have

⁶¹ Kopec (unpublished manuscript), Goldman (2010).

⁶² Elga, (ms.)

⁶³ In terms of degrees of beliefs, taking your standards to be truth conducive means having high confidence that assigning credences using your standards will result in credences with high expected epistemic value, where epistemic value is measured by a scoring rule. A scoring rule is just a function that measures the epistemic "success" of a credence in a proposition. These rules will assign high value to high credences in truths and low credences in falsehoods, and low value to high credences in falsehoods and low credences in truths.

been raised for permissivism, which I think are more serious. These worries all center around the idea that a permissivist is committed to thinking of her beliefs as, in some sense, “arbitrary”.

(1) *The Evidence Pointing Problem (Sosa, White)*

One way to argue against permissivism is to think about where the evidence “points.” The kind of permissivist that I am interested in will sometimes want to say that it is permissible, given E, to believe p and permissible, given E, to believe ~p.⁶⁴ But surely it should only be rationally permissible to have beliefs that the total body of evidence supports, and it is impossible for the evidence to support belief in both p and ~p. After all, to whatever extent the evidence supports believing p, it supports disbelieving ~p, and vice versa.

I think that the imagery being appealed to in this argument misrepresents the permissivist position I am defending. The permissivist I have in mind does not think that there is a special evidence “dial”, which points to the degree of confidence one should have in p, given E, and that, sometimes, the dial points in two directions at once (being confident in p and being confident in ~p). Rather, the permissivist I have in mind thinks that there are multiple evidence “dials” corresponding to different permissible ways of weighing the evidence (different epistemic standards). Sometimes, all of the permissible standards warrant the same attitude, but other times, the different standards warrant different attitudes.

Less metaphorically, the reason it can be permissible for Anna to believe p, and Bob to believe ~p, is *not* that there is a special set of epistemic standards, such that according to these standards, it is reasonable to believe p and reasonable to believe ~p. Rather, the permissivist thinks that what it is reasonable to believe about p needs to be understood relative to some set of epistemic standards. Thus, the reason it is permissible for Anna and Bob to differ in their beliefs with regard to p is that Anna and Bob have *different* sets of standards, which differ regarding whether to believe p, given their body of evidence. Crucially, no one set of epistemic standards will ever warrant belief in p and ~p.

⁶⁴ See footnote 52 for an explanation of why I am restricting my discussion of permissivism to views according to which it is sometimes permissible to believe p and permissible to believe ~p given E.

(2) *A Cluster of Worries: Arbitrariness (Christensen, Feldman, White)*

The general worry raised by arbitrariness is as follows: if you think that, given E, it is reasonable to believe p and reasonable to believe ~p, it looks like having one, rather than the other, belief is *arbitrary*. For if the evidence supports believing p to the same extent that it supports believing ~p, what reason could you have to prefer your particular position?

Roger White uses an example which makes this worry especially pressing. He has us imagine a case in which the permissivist wants to say that, given E, it is permissible to believe p, and, given E, it is permissible to believe ~p. We are then asked to suppose that there are two pills with the following feature: if we take pill #1, we will end up believing p, and if we take pill #2, we will end up believing ~p. We are faced with a choice: Either, we could look at some evidence relevant to whether p and come to whatever conclusion we come to on the basis of that evidence, or, we could take a randomly selected belief pill and end up with our belief that way. It seems like it is better, from the point of view of obtaining true beliefs, to arrive at beliefs by a careful examination of the evidence rather than by a randomly selected belief pill. But if one is a permissivist, it seems like one should think that both methods are equally likely to lead to truth – for, after all, both methods are equally rational.

According to the arbitrariness worries, if Sally believes p, given E, while knowing that it is just as reasonable to believe ~p, given E, she should think to herself: “why, given E, should I believe p rather than ~p when both are equally rational?” There are different ways of understanding this worry. On one way of understanding it, it is just a challenge to Sally to justify her believing p, rather than ~p. If this is all the worry amounts to, there is no problem. Sally can justify her belief in the ordinary way: by appealing to evidence, arguments, reasoning, and so forth. For example, maybe Sally can defend her atheism by appealing to arguments like the Problem of Evil, or appeals to ontological simplicity. She can, at the same time, recognize that there are responses to these arguments that can seem compelling, and other arguments in favor of theism (like, perhaps, the Argument from Design). But despite realizing that there are alternative reasonable ways to weigh the evidence, Sally thinks that she is more likely to end up with a true belief by using her own standards of reasoning. (Recall that part of what it is to have standards of reasoning is to take them to be truth conducive.) So Sally will not regard her atheism as arbitrary at all. Atheism is the belief warranted by Sally’s standards, which she takes to be more truth conducive than standards which warrant belief in theism, or agnosticism.

Appealing to an agent's standards of reasoning is also how we can respond to the case that White describes. Let's begin by thinking about how the pill in White's case is supposed to work. There are two possible mechanisms we should consider: The first mechanism would involve the pill changing the agent's belief *without* changing her standards. The second mechanism would involve the pill changing her belief *by* changing her standards. In either case, I will argue, the reason that an agent will prefer to form a belief about whether *p* by looking at the evidence, rather than taking a pill, is that the agent will think that evaluating the evidence is more likely to result in her having a true belief about whether *p*. Let's consider each mechanism in turn, beginning with the first. Why would a permissivist prefer to look at the evidence instead of taking a pill which will cause her to form a belief about whether *p* without changing her standards of reasoning? Because if she takes a randomly selected pill she may end up having beliefs that conflict with her standards of reasoning. Since she thinks that her standards of reasoning are truth conducive, she thinks she is more likely to end up with a true belief by applying her standards than by taking a pill. Let's now consider the second mechanism. This pill would work, *not* by potentially causing the agent to have a belief that conflicts with her standards of reasoning, but by changing the standards of reasoning themselves. This sort of pill will look equally unattractive to a permissivist. Although she knows that, later, she will not be violating her own standards (since she will have new standards), she does not now think that her later standards will be as likely to lead her to a true belief as her current ones. For this reason, an agent concerned with the truth about whether *p*, would refuse to take a pill that would change her standards of reasoning.^{65,66} Thus, although from some neutral standpoint, taking a pill and evaluating the evidence may look like equally good options, they certainly will not look equally good from the agent's own standpoint.

The defender of UNIQUENESS might think that appealing to standards in justifying having one, rather than another, permissible belief just pushes back the question: for what reason does Sally have for thinking that she is more likely to end up with a true belief by weighing the evidence *her way*, using *her* standards, rather than in some alternative way, if both ways are rationally permissible?

⁶⁵ White describes another case which is meant to be problematic for the permissivist and avoids the response I have given here. However, I will postpone discussion of this case to later in the paper where it becomes especially relevant.

⁶⁶ For other responses to White's argument see Kopec (unpublished manuscript), Douven (2009) and Ballantyne and Coffman (forthcoming).

The answer is that Sally has no reason *independent of her standards of reasoning* for thinking that her standards are more likely to lead to the truth than some alternative. But note that the defender of UNIQUENESS is in no better position, for the defender of UNIQUENESS thinks that one set of standards is the uniquely rational set⁶⁷ and if we ask the defender of UNIQUENESS to give us a reason to think that *her* standards are the uniquely rational standards, she cannot do so in a way that is independent of these standards either. Whether we are permissivists or not, we can never give reasons for why we weigh the evidence in one way rather than another that are independent of everything else. This is just a fact about epistemic life that we have to live with: the methods that we use to evaluate evidence are not the sorts of things we can give independent justification for.⁶⁸

Here is another way of seeing what is going on: there is a sense in which Sally thinks of alternative standards as “just as good” as her own and a sense in which she does not. For Sally thinks that although her standards are more truth conducive than some alternative, other standards may be just as rational. We might cash this out by thinking that the principles of rationality are going to be *general*: they will be principles about what kinds of considerations count in favor of what kinds of hypotheses. But these sorts of general considerations are not sufficiently robust to pin down a unique doxastic state given by any body of evidence. So even if Sally and her friend both conform to these principles, their standards may differ with regard to

⁶⁷ A set of standards is rational just in case following these standards will result in rational beliefs.

⁶⁸ You might think that although the non-permissivist and the permissivist are equally unable to justify their set of standards in a way that is independent of the standards themselves, the non-permissivist can do *something* that the permissivist cannot do. For the non-permissivist, perhaps, can justify the rationality of some *particular* standard by an inductive inference from the rationality of other standards. For example, suppose that the non-permissivist in question is an objective Bayesian. She thinks that the unique probability to assign towards the proposition that there exist pink elephants, (call this “p”), given her evidence, is .0023. Perhaps she can think to herself: “I have lots of other rational standards, which are best explained by my having some excellent rational capacities, so it is likely that this standard is rational as well.” Although, like the non-permissivist, the permissivist can justify the *rationality* of some standard by appealing to the rationality of others, she cannot always infer from the fact that a bunch of her standards are *truth conducive* to the claim that some other cluster of standards is truth conducive. This is because, on the permissivist’s view there, is no *capacity* which explains why, of the variety of rational standards she might adopt, she has adopted standards that are truth conducive. The permissivist will think that what led her to adopt standards that are truth conducive rather than some alternative set of rational standards is sheer luck. For this reason, she cannot infer future successes from past successes. But far from being an objection, the claim that there is no capacity which can explain why an agent’s standards are truth conducive, is at the very heart of what the permissivist wants to assert! As mentioned earlier, part of the motivation for permissivism is the thought that it is implausible that there is some kind of capacity which will determine a unique appropriate attitude to adopt given any body of evidence. Rationality can only narrow down the set of possible standards, so far. Beyond that, the people that end up with standards that lead them to have more true beliefs are just lucky. So although there may be *some* sense in which the non-permissivist can justify her standards and the permissivist cannot, the possibility of such a justification is exactly what the permissivist *wants* to deny. It is therefore dialectically ineffective to point out that the permissivist cannot justify her standards in this particular way as an argument against permissivism.

how exactly they weigh the different considerations and thus, in any given case, Sally and her friend might rationally come to different conclusions.

I think that the problem with this cluster of arguments for UNIQUENESS is that they all, in some way, rely on one of two false assumptions. Either they assume that the permissivist cannot justify her belief in permissive cases, or they assume that our fundamental standards of reasoning need to be justified independently of those standards themselves. The first assumption is false because a permissivist can always make a case for her belief in the usual way. The second assumption is false because regardless of whether you are a permissivist, a justification for our standards of reasoning is not something we can provide independent justification for and the demand for such justification would result in widespread skepticism.⁶⁹

3. How Permissivism Bears on Irrelevant Factor Cases (Defense of P2)

In the previous section, I provided some motivations for thinking that, at least in some cases, there is more than one way to rationally respond to a particular body of evidence. From this point onwards, I will be assuming that permissivism is true. In this section, I show how permissivism bears on cases in which we learn that our beliefs were influenced by an irrelevant factor. In fact, the connections between the permissivism debate and the question about irrelevant influences on belief are quite intricate. Some of the arguments leveled against permissivism deal with cases of irrelevant influences (one such case is the case of the pills, and there is more to come), and some judgments on how to deal with irrelevant influences make assumptions about permissivism. The goal of this section is to argue for the second premise of the main argument. That is, that given permissivism, many cases of irrelevant influences need not worry us. This is because many of the cases in which we worry about the irrelevant influences on belief are *permissive*, and, as I will show, in permissive cases, learning about an irrelevant influence doesn't justify a reduction of confidence.

Let me begin by introducing you to my opponent. My opponent is someone who thinks that learning that our belief was caused by an irrelevant influence should always cause us to decrease confidence in that belief, even in *permissive* cases.⁷⁰ Here is how I will proceed: First, I will try to pinpoint what is worrying about irrelevant influence cases (IF cases) and come up with

⁶⁹ Adam Elga has defended this point in his paper "Lucky to Be Rational" (ms.).

⁷⁰ Ekaterina Vavova (ms.) has explicitly endorsed this view, but it is implicit in many of the arguments that favor reducing confidence in IF cases.

two hypotheses which might explain the worry. The first hypothesis, I argue, is true, but yields the result that we do not need to worry about permissive IF cases. The second hypothesis gives the result my opponent wants but is inconsistent with permissivism, and hence, false. Thus we arrive at the happy conclusion that, in permissive cases, we do not have to worry about irrelevant influences on our belief.

3.1. What's Wrong with Irrelevant Influences: Some Easy Cases and Two Proposals

I will start with a discussion of what we might find disconcerting about IF cases in the first place. To warm up, let's start with a very easy case.⁷¹

DRUG: You have worked through a logic problem and concluded that *p*. You then learn that your evil logic teacher flipped a coin. If the coin landed heads she did nothing but if it landed tails she slipped a reason-distorting drug in your coffee. People who reason through logic problems under the influence of this drug nearly always get the answer wrong. You wonder whether you ought to reduce confidence in *p*.

In DRUG, it seems clear that you should decrease confidence in *p*.⁷² There are two rough candidate explanations for why this is the case:

The “My belief might not be rational!” hypothesis: You should reduce confidence because you realize, upon learning about the drug, that there is a significant chance that the belief you ended up with is *irrational*.⁷³

The “My belief might not be true!” hypothesis: You should reduce confidence because you realize, upon learning about the drug, that there's a significant chance that the belief you ended up with is *false*.

These two explanations have not been distinguished in the literature, but I think the difference between these two explanations is important, *especially* if one is a permissivist. So

⁷¹ This case is based on a case described by David Christensen (2010).

⁷² Weatherson (ms.) and perhaps also White (2010)^b think that, even in a case like DRUG, you need not reduce confidence upon learning about the irrelevant influence. If they are right, then one *never* needs to worry about being in an IF case, permissive or not. What I am aiming to do in this paper, however, is show that we can maintain the intuitive judgment that, in cases like DRUG, we *should* reduce confidence, without thinking that we need to reduce confidence in all cases of irrelevant influences. So for the remainder of the paper I will be setting aside the view that learning that one is in an IF case never gives one reason to reduce confidence.

⁷³ Here, and elsewhere where I discuss this hypothesis, it is important to note that the worry raised is one about the rationality of the belief *prior* to learning about the irrelevant influence. More precisely, if *E* is your evidence before you learn about the drug, and *E** is your evidence after you learn about the drug, this hypothesis says that the reason to be worried is that, upon learning about the drug, you realize that *E* might not support your belief.

the task ahead is to figure out which of these two (rough) explanations is the correct one. However, before moving forward, we need to make some refinements to our hypotheses. Here is why: Suppose that you are in DRUG and I tell you that you should reduce confidence; either because you have reason to suspect that your belief is irrational, or because you have reason to suspect that your belief is false. You might respond as follows: “Look, I grant that *before* solving the logic problem, learning about the drug would give me reason to suspect that I would come to have an irrational or false belief. But now that I have reasoned through the problem I know that, in fact, I did not reason poorly. For I now know that the correct answer to the problem is *p*, and I know that *p* is entailed by my evidence, and furthermore, I believe *p*! So I have arrived at the correct and rational answer. I guess I didn’t take the drug after all!” If you think that you should reduce confidence in DRUG, you will not be impressed with such reasoning. For you will think that when determining how likely you are to reason poorly, or make a mistake, you need to do so in a way that is *independent* of your reasoning about *p*. In other words, you cannot appeal to the very reasoning which is in question to defend the claim that you were not under the influence of the drug when reasoning. Frequently, to figure out how likely we are to be rational or correct in a way that is independent of the reasoning in question, we can do so by thinking about how we would have judged our chance of success *before* doing the reasoning. Here, then, are the refined versions of our hypotheses. I will call the refined versions of the two hypotheses mentioned above “RATIONAL INDEPENDENCE” and “TRUTH INDEPENDENCE.”⁷⁴

RATIONAL INDEPENDENCE: Suppose that *independently of your reasoning about p*, you reasonably think the following: “were I to reason to the conclusion that *p* in my present circumstances, there is a significant chance my belief would not be *rational!*” Then, if you find yourself believing *p* on the basis of your reasoning, you should significantly reduce confidence in that belief.

TRUTH INDEPENDENCE: Suppose that *independently of your reasoning about p*, you reasonably think the following: “were I to reason to the conclusion that *p* in my present

⁷⁴ Principles along the lines of RATIONAL INDEPENDENCE and TRUTH INDEPENDENCE have been suggested by David Christensen(2007), Adam Elga (2007) and Ekaterina Vavova (ms). Although there are known difficulties with extracting a portion of our beliefs, the claim that *sometimes* it is unclear how exactly to set aside one’s reasoning about *p* does not make such principles useless. In fact, in many cases, it is quite straightforward how to set aside one’s reasoning about *p*, especially when one can simply refer to one’s attitude at an earlier time (before the reasoning took place). Since in all of the cases that are relevant to the purpose of this paper there are no special difficulties that arise in setting aside one’s reasoning about *p*, it is not necessary to appeal to some general solution to the extraction problem. We could just as well use a version of this principle which was restricted to those cases in which it is possible to set aside one’s reasoning about *p*.

circumstances, there is a significant chance my belief would not be *true!*” Then, if you find yourself believing *p* on the basis of your reasoning, you should significantly reduce confidence in that belief.^{75,76}

In the next section, I will talk about which of these two principles we should accept. But first, I want to describe one more case which will be helpful to keep in mind when thinking about the more realistic cases (like religious belief) that I will be discussing in the next section.

ZAPPER: You start out having no idea whether *p*. You then look at some evidence which either entails *p* or entails $\sim p$. (Suppose that there is only one rational response to this evidence: if it entails *p*, you should believe *p*, and if it entails $\sim p$, you should believe $\sim p$). You reason and conclude that *p*. You then learn that an evil scientist flipped a coin to determine what you will believe. She decided that, if the coin lands heads, she will zap your brain so that it seems to you that there is a good argument for *p*, and if it lands tails, she will zap your brain so that it seems to you that there is a good argument for $\sim p$.

There are some differences between DRUG and ZAPPER, but despite these differences, it should be clear that the same reasons that motivate decreasing confidence in DRUG also motivate decreasing confidence in ZAPPER. For just as in DRUG, it is true in ZAPPER that, if you set aside your reasoning about *p*, it is quite likely that, by reasoning, you would end up with an irrational and false belief. To see why, imagine, once again, that you knew about the zapping set up *before* you formed the belief. At this point, you have no idea whether *p*, and you have no idea whether the evidence you will get supports *p*. You do, however, know that, no matter what, you will end up believing *p*. Thus, you will think that there is a significant chance that your belief in *p* will be false and irrational, and so, no matter which principle is correct, you should decrease confidence in ZAPPER as well.

⁷⁵ Two notes about these principles: first, although I stated them in terms of belief, for simplicity, at least the first principle could be translated straightforwardly into a principle in terms degrees of belief. Second, there is a worry that both of these principles lead to widespread skepticism. To avoid skeptical results, some bells and whistles need to be added which are discussed in Christensen (2007), Elga(2007) and Vavova (ms). Since, for my purposes, it does not matter whether we use these principles, or the fancier ones, I am going to use the simple principles for convenience, though nothing about what follows rests on the particular version of the principle I will be using.

⁷⁶ There is another independence principle which that has been suggested by Adam Elga (ms). Elga’s principle says, roughly, this: if *S* rationally has doxastic attitude *A* towards *p*, it must be true that, independently of *S*’s reasoning about *p*, *S* judges that *A* accords with her standards of reasoning. While this proposal bears some similarities to RATIONAL INDEPENDENCE, I have argued elsewhere (reference omitted for blind review) that this principle provides neither a necessary nor sufficient condition for when you should decrease confidence in IF cases. For this reason, I will not be talking much about this principle here.

3.2 The Argument for Maintaining Belief in Permissive Cases

We are now ready to set aside pills and zappers and consider the kind of case we encounter in real life:

COMMUNITY: You have grown up in a religious community and believe in the existence of God. You have been given all sorts of arguments and reasons for this belief which you have thought about at great length. You then learn that you only have the religious beliefs that you do, and only find the reasoning that you engaged in convincing, because of the influence of this community. If you had grown up elsewhere, you would have, on the basis of the same body of evidence, rejected those arguments and become an atheist.

I am going to assume that COMMUNITY is a permissive case. This is a substantial assumption, which I believe to be correct, but I will not be arguing for that here. If you do not think the case is permissive, just plug in your own favorite permissive case. Now the question is as follows: do the hypotheses we've considered that motivate decreasing confidence in DRUG and ZAPPER also motivate decreasing confidence in COMMUNITY? Note that, structurally, COMMUNITY is very much like ZAPPER. For, in effect, your community acts like the evil scientist. Your community determines what belief you will come to have about religious matters. Ultimately, I will argue that you need not significantly decrease confidence in COMMUNITY by arguing that, in general, you should not decrease confidence in permissive IF cases.

The general point I will be making is as follows: if you start out with E, and, you know that on the basis of E, it is permissible to believe p and permissible to believe \sim p, learning that your belief in p was caused by an irrelevant factor does not give you reason to decrease your confidence in p. The argument for P2 of the main argument will proceed as follows:

Argument for P2

P3. The view that you should reduce confidence in permissive irrelevant influence cases is motivated by TRUTH INDEPENDENCE.

P4. TRUTH INDEPENDENCE says to decrease confidence in all permissive cases (even when there are no irrelevant influences!).

P5. If you have to give up your belief in all permissive cases, there are no permissive cases. (definition of permissivism)

P6. TRUTH INDEPENDENCE is inconsistent with permissivism. (P4, P5)

C2. If permissivism is true, the view that you should reduce confidence in permissive irrelevant influence cases is unmotivated. (P3, P6)

I will first argue for P3, and then for P4.

Defense of P3:

To see why P3 is true, let's continue to use COMMUNITY as our toy permissive case. I will now argue that if *RATIONAL INDEPENDENCE* is the right principle, it is permissible to maintain belief, whereas if *TRUTH INDEPENDENCE* is right, you must give up your belief in COMMUNITY. Recall that *RATIONAL INDEPENDENCE* says to decrease confidence if you worry that your belief in p might be irrational. In the case we are imagining, however, the community not only caused you to believe in God, but instilled in you rational standards of reasoning that warrant belief in God (this is why you find the arguments for theism plausible). If you had grown up in a different community, you would have been instilled with a different set of rational standards which would have warranted atheism. Since the case is permissive, you were guaranteed to end up with a rational belief *no matter what*. So if worries about the *rationality* of one's beliefs are what explain why sometimes we should decrease confidence in IF cases, we will never have to be worried by irrelevant influences in permissive cases.

What about *TRUTH INDEPENDENCE*? If *TRUTH INDEPENDENCE* is the correct explanation of why we should sometimes decrease confidence in IF cases then you *should* decrease confidence in COMMUNITY. Why? Because *independently of your reasoning about religious matters*, you should think that there is a significant likelihood that you would form a false belief as a result of your community's influence. (As in ZAPPER and DRUG, you can see this by imagining that you are judging the likelihood of forming a true belief about God's existence *before* you joined the community and reasoned about the question. In this state you know that you will come to believe in God on the basis of your community's influence, but you are agnostic about God's existence. Thus, you will think that there is a significant chance that the belief you will form will be false. This is the perspective you have to occupy when you set aside your reasoning about religious matters).

Thus, whether we should decrease confidence in COMMUNITY boils down to which of our hypotheses is the correct one. In the next subsection I will argue that we should reject *TRUTH INDEPENDENCE* and instead explain the easy cases we started out with using *RATIONAL INDEPENDENCE*. Since, in COMMUNITY, you can be confident that the belief you formed was rational, you do not need to abandon your belief.

Defense of P4:

Here is where we are: I have suggested two principles which might explain why we should sometimes decrease confidence when learning that our beliefs were caused by irrelevant factors. As a reminder, here are the two principles once again:

RATIONAL INDEPENDENCE: Suppose that *independently of your reasoning about p*, you reasonably think the following: “were I to reason to the conclusion that p in my present circumstances, there is a significant chance my belief would not be *rational!*” Then, if you find yourself believing p on the basis of your reasoning, you should significantly reduce confidence in that belief.

TRUTH INDEPENDENCE: Suppose that *independently of your reasoning about p*, you reasonably think the following: “were I to reason to the conclusion that p in my present circumstances, there is a significant chance my belief would not be *true!*” Then, if you find yourself believing p on the basis of your reasoning, you should significantly reduce confidence in that belief.

I have argued that the worry about irrelevant influences in cases like COMMUNITY (where you learn that your religious beliefs were caused by the community you grew up in) arises only if TRUTH INDEPENDENCE is correct. In this subsection, I will argue that we should reject TRUTH INDEPENDENCE, and hence, we can rationally maintain belief in such cases.

Here is how the argument will proceed. My opponent wants a principle like TRUTH INDEPENDENCE to explain why you should give up belief in cases of irrelevant influences. This means that she needs a hypothesis which will distinguish a case in which you do some reasoning and believe in God on its basis *without* any irrelevant influences, from a case in which you do some reasoning and believe in God, but as a result of an irrelevant influence (as in COMMUNITY). I will grant to my opponent that TRUTH INDEPENDENCE yields the result that one should decrease confidence in COMMUNITY. However, I will argue that TRUTH INDEPENDENCE will tell you to give up your belief in permissive cases *all the time*, even when everything is hunky dory. So, as a matter of fact, if TRUTH INDEPENDENCE is correct, *there are no permissive cases*. Any time you try to have a belief in a permissive case, TRUTH INDEPENDENCE will tell you to give it up. Thus, the principle will fail to distinguish permissive cases in which there are irrelevant influences, from cases in which there are not.

To see why TRUTH INDEPENDENCE rules out any permissive cases, consider a case in which an agent reasons just as you do in COMMUNITY, but this time, without any irrelevant influences. We can make this the most idyllic permissive case imaginable. This agent,

Caveman, is a perfectly rational being who sits in a cave and carefully considers the arguments for and against the existence of God. No drugs, no zaps, no communities. Just pure, unadulterated, reason. Suppose that after careful deliberation, he comes to believe that God exists (let's call this proposition "G") on the basis of the same reasons that you are given in COMMUNITY. He also recognizes that someone could rationally reject G on the basis of the same evidence that he has. If permissivism is true, cases like this must exist.⁷⁷ But, as I will show, even in this idyllic case, TRUTH INDEPENDENCE will tell the agent to give up his belief in G.

For Caveman to determine whether he can maintain his belief in God given TRUTH INDEPENDENCE, he must think about how likely he is to be right, in a way that is independent of his reasoning about the existence of God. As in the other cases, we can think of the perspective from which Caveman has to reason as the one that Caveman was in before he went into the cave to deliberate about the existence of God. From this point of view, how likely should he think it is that, were he to reason to the conclusion that God exists, this belief would be true? Answer: not very likely.

The reason is as follows: once Caveman has set aside his reasoning about the existence of God, what he is left with is a perspective that is *neutral* with respect to the existence of God.⁷⁸ *From this neutral state*, he should think it is unlikely that he will form a true belief about God's existence on the basis of his reasoning. This is because he knows that the case is a permissive case, and so, both G and ~G are rational given his evidence. If he cannot appeal to the actual reasoning and standards that led him to the conclusion that he reached, he has no reason to expect that, of the two rational positions he might adopt, he will come to have the correct one. Thus, even in this idyllic permissive case, TRUTH INDEPENDENCE would tell Caveman to give up his belief in G. And so, if we accept TRUTH INDEPENDENCE, we get the result that *even when there are no irrelevant influences*, we should give up belief in permissive cases. This means that

⁷⁷ You may be able to get around this argument if you think that, although permissivism is true, one can never know that one is in a permissive case. This, however, is a very implausible view and accepting it requires giving up on some of the considerations that motivate permissivism in the first place.

⁷⁸ Since the reasoning, in this case, is the same as the reasoning in COMMUNITY, my opponent must grant that setting aside the relevant reasoning, in this case, results in a neutral perspective to get the result that setting aside one's reasoning, in COMMUNITY, results in a neutral perspective. Note that part of what will have to be set aside in permissive cases like COMMUNITY are the standards of reasoning that are in question. (Since we are assuming that, in permissive cases, what explains disagreement is a difference in standards, obtaining a neutral perspective will require setting aside the contentious standards).

there could be no idyllic permissive cases like the ones I described. Since permissivism says that there could be such cases, TRUTH INDEPENDENCE is inconsistent with permissivism.

Here is a somewhat deeper explanation of what is going on: If you are a permissivist, setting aside your reasoning (in permissive cases) will leave you in dire straits. Recall that the permissivist was faced with an arbitrariness worry. Why, if it is rational to believe p and rational to believe $\sim p$, should you think that *your* belief (whichever one it is) is the right one? I argued that the permissivist should deny the accusation that her belief in p is arbitrary. The reason, I claimed, the permissivist should not regard her belief as arbitrary is that she can appeal to her particular standards, reasons, and arguments for believing p . But TRUTH INDEPENDENCE says that for a belief to be rational, we must think it likely to be true, *even if we set all of that aside*.

Note that if you are not a permissivist, the condition set by TRUTH INDEPENDENCE will frequently be met. For even if you set aside your actual reasoning, when there is no funny business going on, it is safe to assume that you reasoned *rationally*. Since rational beliefs tend to be true, you do not need to appeal to any actual reasoning you did to justify the claim that the belief you ended up with will likely be correct. In permissive cases, however, this is not an option. For if believing p and believing $\sim p$ are both rational responses to a body of evidence, we cannot think it likely that our belief in p is true on the basis of it being rational, since the very same reasoning would lead us to think that a $\sim p$ belief would be true.

In sum, if, in a permissive case, we must judge it likely that we would come to the right conclusion, even setting aside our reasoning about the issue, the arbitrariness worry returns, and this time, with the upper hand. For once you have set aside the actual reasoning that led you to p , you have no reason to think that the belief you end up with will be true. In other words, what TRUTH INDEPENDENCE demands is exactly what the permissivist cannot provide: an *independent* reason for thinking it likely that her beliefs, in permissive cases, are true. But this has nothing to do with irrelevant influences. The permissivist can never provide this kind of independent justification, and should reject any principle which suggests that such justification is required.

3.3. An Application

Here is what happened so far: I have argued for permissivism, and I have argued that, if permissivism is true, we need not significantly decrease our confidence in permissive IF cases. In the next section I am going to present a problem that arises for my view, but before doing so I

want to consider one more case that has been much discussed in the literature. This case was first described (autobiographically) by G.E Cohen (2000):

SCHOOL: You are a graduate student in philosophy at Oxford and come to believe that there exists an analytic/synthetic distinction. You then realize that if you had gone to graduate school at Harvard you would have come to reject the existence of the analytic/synthetic distinction.

Assuming that the question of whether or not there is an analytic/synthetic distinction is (or was) a permissive one, this case may look quite similar to COMMUNITY, and one might have thought that, for this reason, this is a case in which it is permissible to maintain one's own belief. However, I think that there is a distinction between these cases which may turn out to be important: namely that, in COMMUNITY, the relevant beliefs and standards of reasoning that you formed, were formed early in life, while in SCHOOL, the relevant beliefs and standards of reasoning that you formed were formed later in life.

Why should the time at which you formed the belief make a difference? Because I think it is not implausible that, in SCHOOL, there is a fact about what belief my standards sanction regarding the analytic/synthetic distinction prior to my having gone to graduate school. If this is true, then learning that I have adopted the beliefs that I have because of the school I went to gives me a reason to be concerned. This is because this realization gives me a reason to think that I may have not *rationaly* come to the belief that I currently possess. This belief may, in ways unbeknownst to me, be in tension with other beliefs and standards of reasoning that I have, and so the belief may be irrational. As I mentioned earlier, I think that we have good reason to decrease confidence in IF cases when there is a concern that our belief is irrational (by RATIONAL INDEPENDENCE), and SCHOOL might be such a case.

The earlier on my belief in *p* was formed, the less likely it is that this belief is in tension with other beliefs and standards of reasoning that I possess. This is because it is likely that many of the beliefs and standards we formed early on are deeply entrenched in the way we think about a host of other issues. Since the beliefs we formed early on had a strong influence on the construction of our (hopefully) coherent system of beliefs, it is unlikely that these beliefs are in conflict with the rest. However, if we form a belief later in life, and learn that it was caused by an irrelevant influence, we might have reason to think that, independently of the reasoning in question, it is likely that we've reasoned *irrationally*, by reasoning in ways that are inconsistent

with other beliefs and standards that we have. For this reason, I am more certain that in cases like COMMUNITY we are permitted to maintain belief even upon learning about the irrelevant influences on that belief, than I am about the permissibility of maintaining confidence in cases like SCHOOL.

4. A Problem

So far I have argued that, while in cases like DRUG, which are non-permissive, we should reduce confidence in our beliefs (because of RATIONAL INDEPENDENCE), the considerations which motivate reducing confidence in DRUG do not give us reason to reduce confidence in permissive IF cases. Given my defense of permissivism, this may seem surprising. For recall that, in defending permissivism against arbitrariness worries, I argued that a permissivist should not be willing to take a pill which will cause her to adopt a new set standards, *even if those standards are rational*. And yet here I am suggesting that if you find out that, unbeknownst to you, you *did* take a pill (or get zapped, or grow up in a certain community) which caused you to adopt some rational set of standards, then you have no reason to abandon your belief. In effect, then, what I am suggesting is this: don't take the pill, but if you do, don't worry about having taken it. There is a seeming tension here which I aim to bring out this section.

We can make this problem more precise by noting that the judgment that we need not reduce confidence in permissive cases seemingly conflicts with the following plausible principle:

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

A version of this principle was first introduced and defended by Bas 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 credence in p is, you should adopt that credence now.

To see how a REFLECTION violation can arise, consider a case that Roger White raises in defending UNIQUENESS. The case is very similar to SCHOOL, except in this case, we imagine that you learn about the irrelevant influence *before* attending school.

BEFORE SCHOOL: You are currently agnostic about whether or not there is an analytic/synthetic distinction (p). You are going to study at Harvard or Oxford where you will hear lots of arguments about the distinction. You know that if you go to Harvard you will come to believe p, and that if you go to Oxford you will come to believe \sim p on the basis of the very same evidence. *Furthermore, you know that in either case, your belief will be a rational one.*⁷⁹ The mail comes and you learn that you were accepted to Oxford and rejected from Harvard.

Once you learn that you will be going to Oxford, you learn that you will, in the future, rationally believe that there is an analytic/synthetic distinction. So, according to REFLECTION, you should believe in that distinction now. But it would be crazy go from agnosticism about the analytic/synthetic distinction to belief in it merely by deciding what school to go to!

To solve this problem, I will argue the permissivist has *independent* reason to reject REFLECTION, and so a violation of this principle is not problematic for her. Additionally, the permissivist can accommodate what is plausible about REFLECTION by endorsing a variant of the principle which I will call PERMISSIVE REFLECTION that does not yield the problematic result. To begin, let's see why the REFLECTION principle, *as stated*, is one that the permissivist should reject. Suppose that I am a permissivist and am considering whether to defer to an expert about p. If I were not a permissivist, it seems like all I would need to know to make it the case that I should defer to this expert is that the expert has more evidence than I do, and that the expert always has attitudes that are rational given her evidence. But if I am permissivist, this is not a sufficient reason to defer. In order for it to make sense for me to defer, I must think that, not only are the expert's attitudes always supported by the evidence, *but that the expert weighs evidence in the same way I do*. To see why, consider the following example: Suppose that I am an atheist and think that theism and atheism are equally rational. I may think that my neighborhood priest knows much more about the bible than I do, but I will not defer to him about whether or not Jesus performed miracles because I think that he has different standards of reasoning than I do. In this case, I think I may be more likely to believe the truth by using my

⁷⁹ Although earlier I suggested that cases like this might *not* be permissive since you may, unbeknownst to you, be violating your own standards of reasoning, I will be supposing here that this *is* a permissive case. I can assume this unproblematically since this case is supposed to raise a problem for my view, and the problem requires cases like this to exist.

own standards with less information, than by deferring to the priest who has more information, but adopts different standards than my own.

Recall that one of the primary motivations for REFLECTION is that we should consider our more informed future selves to be experts and, hence, should defer to them for the same reason we defer to experts. Therefore, just as a permissivist should not always defer to an expert with different standards than her own, she should also not always defer to a future time slice of herself with different standards than her own. The reflection principle, then, that the permissivist should accept is as follows:

PERMISSIVE REFLECTION: If you know that, in the future, you will, without loss of information, rationally have doxastic attitude A towards p, *and your future self has the same standards of reasoning as your current self*, you ought to now have doxastic attitude A towards p.

It is crucial in the case we have been discussing that your future self has *different* standards than your current self. For recall that if the Harvard and Oxford students permissibly have different beliefs about the analytic/synthetic distinction, this difference must be explained by appealing to a difference in their standards. Thus, if you know that you will reasonably believe p in the future, you know that, when you go to school you will adopt some standards that you currently do not possess (since you are now reasonably agnostic about p, we can suppose that your current standards are either silent about whether p, or dictate that you should be agnostic, while your future standards will bring about a belief in either p or \sim p). Since you know that your standards at the later time will differ from your current standards, the version of REFLECTION that the permissivist will accept will not imply that you ought to believe or disbelieve p merely on the basis of deciding what school to go to.

5. Disagreement

I would like to end with a few brief remarks about how what I have said bears on the debate about how to react to peer disagreement. There has been much discussion in recent years about whether or not learning that an epistemic peer disagrees with us, gives us a reason to

suspend judgment in that belief.⁸⁰ I think that much of what I have said here about irrelevant influences has applications to the disagreement debate as well.

The views in the disagreement literature according to which you should decrease confidence in the face of peer disagreement are motivated by the very same kind of independence principles which motivate decreasing confidence in IF cases. The thought is that if somebody is my peer and we disagree, then, independently of the reasoning in question, the probability of my being right is equal to the probability of my peer being right. This means that, independently of the reasoning in question, I can't assign a very high probability to my belief being true.

As we have seen, learning that, independently of your reasoning about *p*, you judge there to be a significant chance that your belief is false, is not enough to motivate decreasing confidence. The correct independence principle is one according to which you should decrease confidence in cases in which, independently of the reasoning in question, you don't assign a high probability to your belief being *rational*. Thus, whether we should decrease confidence in cases of peer disagreement is going to depend on whether or not the case is permissive, just as whether or not we should decrease confidence in IF cases depends on whether or not the case is permissive.

This result can, I think, shed some light on a number of problems that have been discussed in the context of disagreement. According to conciliatory views,⁸¹ you should suspend judgment when you learn about a peer's disagreement. But there is a problem with such views, which Elga (2007) calls the "spinelessness worry." The worry is that if we think we should suspend judgment in cases of peer disagreement, we will have to suspend judgment on a large number of our important beliefs (given how much disagreement there is), and the result will be that that we will be epistemically "spineless."

While the conciliatory views can be accused of, perhaps, too much humility, the alternative views (sometimes called "steadfast views"⁸²) can be accused of too much epistemic arrogance. One of the much discussed cases in the disagreement literature is the following:

⁸⁰ See, for example, Kelly (2005), Elga (2007), Christensen (2007) and Feldman (2007).

⁸¹ Conciliatory views have been defended by Elga (2007), Christensen (2007) and Feldman (2007).

⁸² Such a view was defended in Kelly (2005).

RESTAURANT: You and your friend are dining together at a restaurant. You get the bill and decide to add an extra 20% for tip, and then split the total. You each do the calculation to determine how much you owe and come to different conclusions. This is not the first time this has happened. In fact, since you are frequent dining partners, this has happened a great number of times and in half of those times it turned out that you were right, and in the other half, it turned out that your friend was right.⁸³

It seems obvious that, in a case like this, you should decrease confidence in your answer. Doing otherwise would display an inappropriate kind of epistemic hubris. And yet the steadfast views, according to which peer disagreement does not give you a reason to decrease confidence, say that, assuming you calculated correctly, you should not decrease your confidence.

The view that I am advocating can solve both of these problems. Since whether or not we should give up a belief in the case of peer disagreement depends on whether or not the case is permissive, there will be some cases in which we should decrease confidence and some cases in which we should not. In RESTAURANT, it is quite clear that, independently of the reasoning in question, I have good reason to think I have reasoned irrationally and so I should indeed decrease confidence. But decreasing confidence in cases like this does not lead us to spinelessness. There will be lots of cases in which I can maintain my beliefs in the face of peer disagreement and these will be permissive cases. As long as I have reason to think that both my peer and I are rational, the disagreement between us will not give me a reason to abandon my belief.

6. Conclusion

I have argued that our theory of rationality should be permissive; that is, that sometimes there are multiple rational responses to given a body of evidence. I have also argued that the way we should react to learning that our beliefs were caused by irrelevant influences depends on whether or not the case is a permissive one. If we think that the irrelevant influence may have caused us to reason irrationally, learning about the influence may give us reason to reduce confidence. In contrast, if the irrelevant influence caused us to have one versus another rational belief, learning about the irrelevant influence doesn't give us a special reason to reduce confidence.

As I mentioned at the beginning, many of the beliefs that play very central roles in shaping our lives were caused by irrelevant influences. In this paper I have focused primarily on

⁸³ This case is first described in Christensen (2007).

our beliefs about the existence of God, but there are several beliefs we have that have been caused by the influence of our communities, schools, friends and family. I hope to have convinced you that many of these beliefs are not threatened by learning certain facts about their causal history. While this may be a comforting thought, it is important to remember that the argument I have given only works if we take these central beliefs of ours to be ones in which the alternatives are rationally permissible. So while the position I am defending allows us to be dogmatic, in the sense that we can maintain our beliefs on these matters despite knowledge of the influence of irrelevant factors and peer disagreement, this kind of dogmatism is warranted only by a great deal of epistemic tolerance. We can only maintain our beliefs if we recognize that the alternative beliefs in question are just as rational as our own.

References

- Ballantyne, N and Coffman E.J. (forthcoming) "Uniqueness, Evidence, and Rationality." *Philosopher's Imprint*.
- Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Burks, A. (1953). "The Presupposition Theory of Induction." *Philosophy of Science*, 20 (3): 177-179
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago, IL: Chicago University Press.
- Cohen, G. A. (2000) *If You're an Egalitarian, How Come You're So Rich?* Cambridge, MA Harvard University Press.
- Christensen, D. (2010) "Higher Order Evidence." *Philosophy and Phenomenological Research* 81 (1): 185-215.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2):187-217.
- Douven, I. (2009) "Uniqueness Revisited." *American Philosophical Quarterly*, 46 (4).
- Elga, A. (2007) "Reflection and Disagreement." *Nous*, XLI (3): 478-502.
- Elga, A. (2010) "Subjective Probabilities Should be Sharp." *Philosophers' Imprint* 10 (5).
- Elga, A. (ms.) "Lucky to Be Rational."

- Elgin, C. (1996). *Considered Judgment*. Princeton, NJ: Princeton University Press.
- Feldman, R. (2007) "Reasonable Religious Disagreement" in L. Antony ed. *Philosophers Without God*. Oxford, UK: Oxford University Press. 197-214.
- Goldman, A. (2010). "Epistemic Relativism and Reasonable Disagreement" in R. Feldman and T. Warfield eds., *Disagreement*. New York, NY: Oxford University Press.
- Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press
- Horwich, P. (1982) *Probability and Evidence*. Cambridge, UK: Cambridge University Press.
- Jeffery, R. (1965) *The Logic of Decision*. New York, NY: McGraw-Hill.
- Joyce, J.M (2005) "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19.
- Kelly, T. (2005) "The Epistemic Significance of Disagreement" in T. Szabo Gendler and J. Hawthorne *Oxford Studies in Epistemology Vol. 1*: Oxford, UK: Oxford University Press: 167-196.
- Kopec, M. (ms.) "In Defense of Epistemic Permissivism."
- Levi, I. (1974). "On Indeterminate Probabilities." *Journal of Philosophy* 71.
- Lycan, W. (1988) *Judgment and Justification*. Cambridge, UK: Cambridge University Press
- Quine, W.V.O and Ullian, J.S. (1970) *The Web of Belief*, New York, NY: Random House.
- Rosen, G. (2001). "Nominalism, Naturalism, Epistemic Relativism." *Philosophical Perspectives* 15:69-91.
- Schoenfield, M. (forthcoming) "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences (And Other Imprecise Doxastic Attitudes)." *Philosophical Studies*.
- Sosa, E. (2010) "The Epistemology of Disagreement" in his *Armchair Philosophy*. Princeton, NJ: Princeton University Press.
- Titelbaum, M.G. (2010) "Not Enough There There: Evidence, Reasons and Language Independence." *Philosophical Perspectives* 24: 477-528.
- Van Fraassen, B. (1984) "Belief and the Will." *Journal of Philosophy* 81: 235–256.
- Vavova, K. (ms.) "Irrelevant Influences."
- Weatherson, B. (ms.) "Do Judgments Screen Evidence?"

White, R. (2005) “Epistemic Permissiveness.” *Philosophical Perspectives* 19 (1):445–459.

White, R. (2010)^a “Evidential Symmetry and Mushy Credence” in T. Szabo Gendler and J. Hawthorne eds. *Oxford Studies in Epistemology Vol.3*. Oxford, UK: Oxford University Press

White, R. (2010)^b “You Just Believe that Because...” *Philosophical Perspectives* 24(1): 573–615.