

Epistemic utility and permissivism

The good news and the bad

Epistemisk nytte og permissivisme

Gode og dårlige nyheter

André Eilertsen



Institutt for filosofi og førstesemesterstudier

Universitetet i Bergen

Masteroppgave i FILO350

Vår 2022

Veileder: Ole T. Hjortland

Abstract

A recent debate in epistemology has been on the question of epistemic *permissiveness*: whether, given some body of evidence, different rational agents are permitted to take different doxastic attitudes to some proposition. This thesis approaches the question from two angles. I first approach the literature on its own terms, surveying and evaluating its arguments. Then I draw on formal work from the field of “epistemic utility theory” to derive a particular solution to the problem. I conclude by reflecting on the philosophical consequences of this solution, in light of the arguments from the first part of the thesis.

Sammendrag

En nylig debatt i epistemologi har tatt for seg spørsmålet om hvor tolerant (eng.: *permissive*) rasjonalitet er: hvorvidt rasjonelle agenter med samme bevis om en påstands sannhet kan komme til forskjellige konklusjoner når det gjelder den påstanden. Oppgaven tar for seg spørsmålet fra to vinkler. Først tar jeg for meg litteraturen på dens egne premisser, og gir en oversikt og evaluering av argumenter på både “ja”- og “nei”-siden. Senere trekker jeg på formelle metoder fra såkalt “epistemisk nytteteori” for å gi mitt eget svar på spørsmålet. Jeg konkluderer ved å reflektere over de mer filosofiske konsekvensene av mitt eget svar, i lys av argumentene i første del av oppgaven.

Acknowledgements

I want to thank, first of all, Ole Hjortland, for sparking an interest in the technical side of philosophy with his introductory lectures in logic, and for discussion over the past year on both theoretical question about rationality and more practical matters. Henrik Dalstø, for his feedback on a draft of chapter 3, and Bernt Barkved, for commenting on late drafts of the first two chapters. The people of the fifth floor, for various enjoyable distractions over the past two years.

Contents

Introduction	5
0.1 The problem	5
0.2 The plan	6
0.3 Overview of the chapters	7
1 Epistemic permissivism - the very idea!	9
1.1 Feldman on reasonable disagreement	10
1.2 White's belief-toggling objections	16
1.3 Evidence and evidential support	23
1.3.1 What is evidence?	23
1.3.2 Two-place or three-place?	25
1.3.3 Weisberg: Univocity, not uniqueness	26
2 The permissivists strike back	29
2.1 Schoenfield's standards	29
2.2 Titelbaum and Kopec	33
2.2.1 Bayesian updating in permissive cases	33
2.2.2 Truth-conduciveness and misleading evidence	36
2.2.3 The Reasoning Room	38
2.3 State of play; looking forwards	41
3 The epistemic risk approach	43
3.1 Believing truly, believing falsely	44
3.2 As someone who loves truth	46
3.3 Spelling out the relevance for permissivism	56
3.3.1 Whither evidence?	56
3.3.2 The attitudes at play	58
3.3.3 What precisely is established?	60

3.4	Are there different legitimate attitudes to epistemic risk? . . .	62
3.5	The <i>via risk</i> for credences	65
4	Intrapersonal permissivism, acknowledged permissive cases, and belief toggling	70
4.1	Intrapersonally permissive cases: arbitrary adoption and belief toggling	71
4.2	Belief-toggling <i>via</i> standards	72
4.3	Acknowledged permissive cases: considering others' standards	74
4.4	Some take-aways and methodological considerations	76
5	Conclusions	79
	Bibliography	82

Introduction

0.1 The problem

The topic of the present thesis is the connection between evidence and rational belief. The central question is whether bodies of evidence uniquely determine rationally permitted doxastic attitudes.¹ The problem has been discussed, for the past couple of decades, under the heading of “epistemic permissiveness,” following a (2005) paper of that name by Roger White. White presents, in that paper and subsequent work (2010, 2014), a number of arguments against the idea that evidence is *permissive*, in the sense that two agents could have the same evidence but be rationally permitted to take different doxastic attitudes to some proposition. He argues, that is, for a principle of *Uniqueness*:² any body of evidence permits, rationally speaking, a single unique doxastic attitude towards any proposition.³ Central to the argument is the thought that permissiveness makes it somehow, in some deep and epistemically worrying sense, arbitrary what our beliefs are vis-a-vis our evidence. What good, the worry goes, is rationality if it doesn’t decide for us which doxastic attitude towards some proposition is the right one, given our evidence?

Other writers have, predictably, joined the debate, coming down on both sides of the issue. Besides White, Richard Feldman (e.g. his 2007) is a canonical *impermissivist*, as is David Christensen (e.g. his 2007). Explicit defenders of permissivism include Douven (2009), Schoenfield (e.g. her 2014,

1. It might be put more precisely as a question of whether rational doxastic attitudes *supervene* on bodies of evidence.

2. White takes the term from the at the time unpublished but circulated Feldman (2007).

3. Some work is needed to parse which precise principle is used by any given proponent of uniqueness, as regards questions of scope and such. I will clarify this in chapter one.

2019), and Kopec and Titelbaum (2019, ms.); though, as the former argues, permissivism is plausibly a consequence of (that is, predicted by) many of our best theories of confirmation.

The problem is of interest for several reasons. It intersects in interesting ways with the literature on the relation between our full and partial beliefs; this, indeed, will be a primary concern in the latter parts of the present thesis. But the question of permissivism also, of course, has bearing on rationality as such: is rationality such a thing as permits different agents with the same evidence to draw different conclusions? It is also of concern for the philosophy of science, which of course cares greatly about what evidence is and how it licences drawing conclusions: can different scientists or groups of scientists, equally reasonable, draw different conclusions from the same body of evidence? Yet another interesting touching-point is the recent debate about disagreement; as we will see soon enough, the purported phenomenon of “reasonable disagreement” served to kickstart the literature about permissivism and uniqueness.

0.2 The plan

As a way to get at the problem of epistemic permissivism I will do two things. Firstly, approach the literature on its own terms, engage with the central figures on both sides of the debate, and evaluate the arguments. This will serve as introduction to and background on the problem and the debate; this takes up the first two chapters of the thesis. Secondly, I will look at some recent work in formal epistemology, and draw on a particular model of belief I think can shed some light on (and motivate) permissivism. The final chapter will put the arguments from the first part to bear on the formal results from the second.

The general idea running through the latter half of the thesis springs out of an observation of William James’s, about the twin epistemic goals of believing truly and not believing falsely being separate, or at least in principle separable (James 1897). Our attitudes to the respective weight or importance of each can lead to different epistemic outcomes with regard to agents’ doxastic attitudes. An early and informal application of this idea to permissivism was given in Kelly (2014); more in-depth work, both within and without the context of permissivism, has been done by e.g. Easwaran (2016), Dorst (2019), Pettigrew (ms.). More work is, of course, needed to unpack

the idea: Are the epistemic goals actually separable, and does separating them actually cash out in differing epistemic outcomes? Presumably not all attitudes to epistemic risk are rationally permissible; which, then, are (not)? Do these differing attitudes to risk only manifest in different full doxastic attitudes, or do they sometimes manifest in different credences among agents with the same evidence?

The central goal of the thesis is to argue, first, that uniqueness is insufficiently motivated, and that, second, some interesting and plausible formal work on rational belief entail permissivism. The approach is chosen for a couple of different reasons. Firstly, simply noting *that* these theories have permissivism as a consequence is interesting, especially when the observation is situated within a broader discussion of that phenomenon. Consideration must then be given as to what weight these observations do or should have in the debate over whether permissivism is true. My conclusion on this point will be modest: it is plausible that our best theories about rational belief, and the relation between course- and fine-grained rational belief, entail some form of permissivism; we ought to accept permissivism to exactly the degree it is entailed by these theories and no further. Secondly, it seems to me a more methodologically sound approach, in a way that illustrates something more general about how epistemology might be done: we give a model of a phenomenon (rational belief), see what this model predicts (and these predictions are quite precise) in cases, and make judgements on the appropriateness of these predicted outcomes.

0.3 Overview of the chapters

In **chapter one**, I will begin digging into the literature, looking at the early papers of White (2005) and Feldman (2007) which set off the debate. I examine their motivations for rejecting permissiveness (by way of embracing Uniqueness), and discuss the particular objections they provide. The end of the chapter also brings together some more general thoughts about the notion of evidence at play in the discussion, how to think about evidential support, and briefly considers an argument from Jonathan Weisberg drawing on empirical work on the psychology of evidence-processing.

Chapter two will take up the dialectical counterpoint to Feldman and White, from Miriam Schoenfield, Michael Titelbaum and Matthew Kopec. The former, we will see, presents an argument for permissivism focused on

the notion of epistemic standards, arguing that different agents can reach different conclusions because they have different such standards. The latter present a barrage of defensive arguments, before giving the reader an example of permissive reasoning anyone would find hard to dismiss.

The focus of **chapter three** will be the epistemic risk approach to permissiveness. I begin by introducing William James's famous observation on the epistemic twin goals of believing truly and not believing falsely, and laying out some formal tools for applying James's idea. I then draw on work by Kenny Easwaran and others in epistemic utility theory and present a model of how course- and fine-grained beliefs ought to cohere in rational agents. I argue that this model has consequences for the permissivism debate: it entails that, given that our evidence is encoded in our credences (as on a Bayesian picture), rational belief is permissive in just the way White and Feldman denies. I develop and evaluate further the idea that there are multiple epistemically legitimate ways of valuing truth and disvaluing falsehood, in the precise guise this idea is presented in the model.

Chapter four will draw out some consequences of philosophical importance of the work done in chapter three. We recall the objections from Feldman and White presented in the first chapter and see how the epistemic risk approach handles them. It is not all good news, but I argue that there is nothing disqualifying for the approach in these considerations.

Chapter five concludes the thesis, highlighting the main take-aways from the earlier chapters individually and the thesis as a whole.

Chapter 1

Epistemic permissivism - the very idea!

There are two immediate responses one might have to the idea of epistemic permissivism. One is that this all seems rather common-sensical: rational people disagree all the time.⁴ The other is that this is all obviously mistaken: the evidence is what it is and says what it says, and rationality is response to this and no more. The latter response will be the focus of this chapter.

I begin the chapter by laying out and discussing an early view in the literature, Richard Feldman's objection to (what was not yet known as) permissivism in his 2007 paper "Reasonable Religious Disagreements." Feldman raises a series of concerns about situations in which two agents with the same evidence seem to rationally reach contradictory conclusions on some question. The concerns revolve around the seeming arbitrariness of rational belief vis-a-vis evidence in such (supposed) situations. Those concerns about arbitrariness continue to evolve in Roger White's treatment of the problem in "Epistemic Permissiveness", which takes us away from disagreement as such into more general matters of belief formation. After surveying the arguments from Feldman and White, I take a step back and think a bit more generally about evidence and evidential support, before briefly considering a marginal but interesting intervention into the permissivism debate by Jonathan Weisberg.

4. Anecdotally: this has been the response from non-philosophers and philosophers with interests outside epistemology when I describe the present project.

1.1 Feldman on reasonable disagreement

Richard Feldman introduced the principle of Uniqueness in the presently relevant form in his (2007), where he takes up the problem of so-called “reasonable disagreement.” More precisely, he is concerned with situations in which two agents, who employ the same methods of reasoning and have all the same evidence (at least all the same evidence pertaining to some relevant hypothesis or discussion), disagree about some proposition, and each of them knows this. It is often claimed that this is perfectly above board from an epistemic point of view; Feldman cites a representative passage from Gideon Rosen:

It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence. ... [I]t 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. (Rosen 2001, 71-2)

Now, locutions of the form “it should be obvious” or “it would appear” (not to mention “surely”) are typically to be regarded with suspicion, but the actual point is, I think, quite correct. It at least *appears* to be a fact of epistemic life that such disagreement occurs. Feldman, as we shall see, disagrees on the factual standing of these appearances.

Feldman goes through a number of reasons why one might believe reasonable disagreement is possible, arguing that each is mistaken. Of the reasons he presents, two are particularly interesting for our purposes: a) the idea that epistemic peers can draw different conclusions from the same body of evidence because the evidence legitimately points in different directions, and b) the idea that different epistemic starting points can allow epistemic peers to end up in different doxastic positions while reasoning rationally from the same evidence. I start with the first.

Feldman uses, throughout his paper, a framing example to illustrate various points; I repeat the example here for ease of reference:

[S]uppose a detective has strong evidence incriminating Lefty and also has strong evidence incriminating Righty of the same crime. Assume that the detective knows that only one suspect could be guilty. One might think that since a case could be made for either suspect, the detective could reasonably believe that Lefty

is guilty and Righty is not, but could also reasonably believe that Righty is guilty and Lefty is not. She gets to choose. If anything like this is right, then there can be reasonable disagreements in the intended sense. If there were two detectives with this same evidence, they could reasonably disagree, one believing that Lefty is guilty and the other believing that Righty is guilty. Each could also agree that the other is reasonable in drawing the contrary conclusion. (Feldman 2007, 204)

As an analysis and/or justification of the two detectives' supposed rationality, apparently from something like a permissivist point of view, I agree with Feldman that this is rather weak. For one, it's not clear that the two agents are actually operating with the same body of evidence; the correct description of the case seems to me to be that each has chosen some subset of the total body of evidence as operative, not considering that the remaining evidence plainly militates against that operative body. A body of evidence, taken as a whole, cannot support, let alone *strongly* support, each of two contradictories for some agent.

Feldman's own point regarding this example I find somewhat less convincing. There's a logical problem, he says, with the idea of reasonable disagreement in this case. If detective A concludes that Lefty is guilty, she can infer that Righty is not guilty. If Righty is not guilty, it follows that detective B is wrong. Thus, Feldman: "if she can draw this inference, she cannot also reasonably think that it is reasonable to conclude that Righty is guilty. This combination of beliefs simply does not make sense" (205). This seems to me to beg the question against reasonable disagreement. Whether such a combination of beliefs makes sense is precisely what is at issue; pointing out *that* this is the combination of beliefs does nothing to answer the problem. If permissivism is true, and the case under consideration is permissive, this combination of beliefs could in fact be perfectly all right.

It is from these considerations that Feldman draws the lesson that evidence is unique, that is, that the Uniqueness Thesis is true. His statement of the thesis goes:

(UTF) A body of evidence justifies at most one proposition out of a competing set of propositions (e.g., one theory out of a bunch of exclusive alternatives) and . . . it justifies at most one attitude toward any particular proposition. (205)

It's not obvious how exactly the case of Lefty and Righty supports such a thesis, and it is not clear that the thesis appropriately reflects exactly how evidential support works. More on the latter point later, but for now I'd like to emphasise that *my* reading of the case, two paragraphs ago, is perfectly compatible with a denial of uniqueness. The case of the detectives and Lefty and Righty is epistemically dubious because it mischaracterizes what the evidence is and what it says, not because evidence and rationality are impermissible. As Douven puts it emphatically: "on *no* extant confirmation theory could one plausibly be said to have *strong* evidence for both of a pair of contradictory proposition" (Douven 2009, 355). Yet, as we shall see, plenty of confirmation theories entail a denial of uniqueness.

Some work is needed to parse precisely what a given formulation of the uniqueness thesis says. There are some unfortunate possible ambiguities lurking; we should, on behalf of everyone in the debate, get clear on where and how these might arise, that we can steer clear of them. In stating the principle, special attention needs to be paid along two axes: firstly, what precisely is uniquely determined by a given body of evidence; secondly, the scope of the unique existential claim.⁵ For the second, it should be noted that there are two readings of the claim that "for any body of evidence, there exists a unique rational doxastic attitude an agent can take to any proposition."⁶ The two different readings come out clearly when formalised:

1) $(\forall \text{ evidence } E)(\forall \text{ proposition } H)(\forall \text{ agent } X)(\exists \text{ attitude } A)$
 $([(E \text{ is } X\text{'s total evidence}) \& (X \text{ adopts an attitude other than } A$
 $\text{toward } H)] \rightarrow [X \text{ is not rational with respect to } H]).$

2) $(\forall \text{ evidence } E)(\forall \text{ proposition } H)(\exists \text{ attitude } A)(\forall \text{ agent } X) ([E$
 $\text{is } X\text{'s total evidence}) \& (X \text{ adopts an attitude other than } A \text{ toward}$
 $H)] \rightarrow [X \text{ is not rational with respect to } H]).$

(Titelbaum and Kopec, ms. 7-8)


The scope of the claim that "there exists some attitude such that ..." makes all the difference here. The two placements give rise to either (1) an intrapersonal or an (2) interpersonal reading of the principle. The former says that any person has a *for them* unique rational attitude to any proposition given

5. I draw on Titelbaum and Kopec, ms. especially pp. 2-3 and 7-8 for the taxonomy of principles.

6. An ambiguity in surface form not unlike examples of the "every philosopher loves some linguist" type.

their evidence, but does not rule out, in principle, that different persons are allowed different attitudes even with the same evidence; the latter reading says that given a body of evidence *everyone* should have the same unique attitude to each proposition. The latter is, clearly, quite a bit stronger than (which is to say, it entails but is not entailed by) the former, and one might reasonably think this is where the action is located.

The distinction between the inter- and intrapersonal uniqueness might also be seen by whether they deny the following situations:⁷



On the left we see a supposed intrapersonally permissive case, where the agent, on obtaining some piece of evidence, can go either way doxastically with regard to P . On the right we see a supposed interpersonally permissive case, with two different agents who, on obtaining the same piece of evidence, can go different ways doxastically with regard to P . The proponent of intrapersonal uniqueness denies that the former is possible, while the proponent of interpersonal uniqueness denies that the latter (though by entailment also the former) is possible. Looking back at Feldman's framing example, the proponent of interpersonal uniqueness would deny that one detective could rationally believe Lefty to be guilty while another rationally believes Righty is, while the proponent of intrapersonal uniqueness would deny merely that a single detective could be permitted to believe either.

On the other axis, there's the question of what exactly is uniquely determined (whether that is intra- or interpersonally). Some formulations have it that evidence uniquely supports some *proposition*, others hold that evidence uniquely determines the *attitude* to have towards that proposition. Feldman, we've seen, presents his thesis as a conjunction of the two, but they are obviously logically distinct principles. It is, however, equally obvious that they are mutually supportive, and that the case for the latter would be made

7. The figure is inspired by Weisberg (2020b, 3).

weaker if the former were rejected.

The question “uniqueness of what?” arises all over again when we consider different doxastic states. Take credences, or partial belief. We could ask whether some body of evidence uniquely supports some other proposition *to a certain degree*, or whether a body of evidence determines uniquely, for any other proposition, which precise (or for that matter, which unique imprecise) credence one ought to have in some other proposition.

In the following we will primarily be concerned with interpersonal attitudinal uniqueness of categorical belief, the principle that any body of evidence determines uniquely which full doxastic attitude any agent is required to have to any proposition. The negation of this principle will be what we call permissivism; I take this to be the more philosophically interesting version of the thesis. In the course of defending permissivism, however, we will be forced to deal also with the notion of propositional evidential support appealed to in Feldman’s first conjunct.

With those clarifications on how to read the Uniqueness thesis out of the way, we can specify that Feldman’s version of the thesis (UTF) is to be understood as the conjunction of interpersonal propositional uniqueness and interpersonal attitudinal uniqueness. Both of the conjuncts will be subject of discussion throughout the present chapter, the latter in §1.2 and the former in §1.3.

The second possible avenue for the permissivist Feldman considers is that of different epistemic starting points, “[w]hether these starting points amount to fundamental claims about the world or epistemological principles about how to deal with evidence” (Feldman 2007, 205). These are, of course, two very different things, and treating them as the same idea leads to confusion. We can (though Feldman does not do this) treat the two as representing two different theories about how evidence works.⁸ The former fits nicely into a Bayesian picture of both evidence and rationality: we come to an epistemic situation with a set of priors, and different priors will, in the face of the same evidence, give different verdicts. This could be illustrated with any number of examples of different credal developments under the same evidence.⁹ The

8. Though this is not to imply that these two theories are in contention; work on connecting them, e.g. by presenting or justifying abduction in Bayesian terms (see e.g. Pettigrew (2021) for a discussion on this body of work.), is on-going.

9. Say $c_\alpha^t(P) = 0.6, c_\alpha^t(Q) = 0.6, c_\alpha^t(P|Q) = 0.90$; and $c_\beta^t(P) = 0.25, c_\beta^t(Q) = 0.5, c_\beta^t(P|Q) = 0.4$. Observing that Q at $t + 1$ would give $c_\alpha^{t+1}(P) = 0.9$ and $c_\beta^{t+1}(P) = 0.4$.

latter fits nicely with various abductive theories, where weighting of different evidential factors or theoretical virtues of the hypothesis considered can lead to different conclusions. An example of this within philosophy is the debate between classically and non-classically inclined logical anti-exceptionalists. Timothy Williamson and Graham Priest both consider the choice of the “right” logic to be continuous with scientific methodology. They are both experts, both, we might think, paragons of rational argumentation, and it is fair to assume they have the same relevant evidence; but “although they list the very same criteria for theory choice (e.g. fit with the evidence, explanatory strength, simplicity, unification), their abductive arguments take them to incompatible conclusions” (Hjortland and Antonsen 2016).

Now, Feldman claims that appeals to (what we might call) priors and abductive weighting only pushes the question one step back: what is the reasonable prior, or the reasonable way of weighting evidence abductively? Surely there must be one uniquely rational answer to both of these, which in turn determines a uniquely rational conclusion? He writes that “[o]nce you see that there are these alternative starting points, you need a reason to prefer one over the other” (Feldman 2007, 206). Those reasons, he claims, ought to be subject to exactly the same rational debate as the original question, and there can be no reasonable disagreement on the former any more than the latter. It’s not clear how reasonable this demand for prior justification of the agent’s epistemic starting point is. For one, it seems that we might have non-evidential reasons for weighing evidence the way we do.¹⁰ Consider the difference in how two agents would weigh evidence in favor and against the existence of God. A person born and raised in a religious environment would quite naturally approach the evidence in a quite different way than someone raised in a secular environment. It’s not clear how one could give reasons for or against either of these starting points. Secondly, there’s the question of how to give reasons for our very faculties of reason-giving. It’s not clear what sort of evidential facts could possibly be given in support of this.¹¹

This is not to say that there are no attempts at uniquely determining epistemic starting points. One such suggestion is commonly known as *objective* Bayesianism (as opposed, obviously, to the *subjective* variant). Bayesianism is often held up as the paradigmatic subjective theory of confirmation

10. I don’t mean to imply that these factors are epistemically legitimate, or that a defence of permissivism would rely on these. They are, however, raised in the debate, and we will see them pop up in Schoenfield’s defence of permissivism next chapter.

11. This response mirrors some approaches to combating scepticism about induction.

and evidence, but one variant of the approach takes priors to be objectively uniquely determined. From this single starting point, then, there will be a unique rational doxastic state for any agent given any body of evidence. Various attempts have been made at uncovering this objective prior, going back at least to Carnap's (1950); these attempts constitute, essentially, rational arguments for which epistemic starting point is correct. The most prominent suggestion on this front is the Principle of Indifference,¹² stating that absent further information (which is obviously the situation when "choosing" priors) one ought to distribute probabilities equally among all options (that is, every member of an exhaustive partition with n members is assigned a probability of $1/n$).¹³ Considerations of both space and relevance rule out extensive discussion of the merits of objective Bayesian here; suffice it to say that if the fate of Uniqueness depends on the truth or falsity of objective Bayesianism, it is far from obvious that its proponents will wind up vindicated.¹⁴

We've at this point extracted the essentials, for our purposes, from Feldman's paper. These regard the uniqueness thesis and the problem of justifying epistemic starting points. Feldman goes on to describe how we ought to behave epistemically in the face of apparent reasonable disagreement, and how his considerations ought to play out in public life on matters political and ethical; these are interesting, but not of much concern to us.¹⁵ We turn now to White's (2005) paper on permissivism, which generalizes Feldman's points away from disagreement as such.

1.2 White's belief-toggling objections

Now, Feldman has framed his discussion in terms of explicit and acknowledged disagreement; he has reached the conclusion that if you know a reasonable person with the same evidence disagrees with you, (at least) one of you is mistaken, and you ought to moderate. Does similar reasoning apply

12. Assumed to hold in a certain sort of cases (games of chance and analogous real life situations) by probability theorists and economists for centuries.

13. White (2010a) presents an argument for a Uniqueness-friendly principle of indifference; Meacham (2013) argues, persuasively, that this fails. Pettigrew (ms.) argues extensively for a wide range of permissible priors, on similar grounds to the Jamesian permissivist view we will discuss in chapter 3.

14. See e.g. (Lyon 2010) for an overview.

15. The upshot is that in many such cases we should take disagreement with rational peers as exerting pressure to moderate our own view.

in cases where there is no explicit disagreement? Roger White, in his (2005), examines such cases and concludes that it does. White presents a number of cases which would assumably be permitted by a permissive view of rationality, and observes that there is something intuitively epistemically wrong, because arbitrary, going on.

Two sorts of cases are relevant, both of which take us away from the epistemology of disagreement as such, and which give rise to questions of inter- and intrapersonal permissivism, respectively. Firstly, the case which is identical to Feldman's except the agents do not communicate; they have the same evidence and reason by the same methods to different conclusions, but because they do not know of each other, there's no pressure to follow Feldman's advice to moderate. A second sort of case involves agents who know or believe justifiably that if they differed in some non-epistemic way they would be drawing different conclusions based on the same evidence and methods of reasoning. White focuses on the latter, taking them to have some bearing on the former. White's formulation of the uniqueness thesis he defends goes:

(UTW) Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition. (White 2005, p. 445)

The formulation is, unfortunately, ambiguous between the intra- and interpersonal readings I distinguished in the discussion of Feldman. White clarifies in his (2014, 312) that his intention is for the *inter*personal reading.¹⁶ We thus have something very close to the second conjunct in Feldman's formulation of the thesis.¹⁷

Failure to disambiguate has consequences, of course; White gives some quite fleshed out examples of cases meant to fend off intrapersonal permissivism, but with dubious force as against interpersonal variants.¹⁸ There is

16. The formulation there goes: "If an agent whose total evidence is E is fully rational in taking doxastic attitude D to P, then, necessarily, any subject with total evidence E who takes a different attitude to P is less than fully rational" (White 2014, 312).

17. Though Titelbaum and Kopec (2019, 206) refer to as (UTW) "personal uniqueness," as distinguished from their designation of (UTF)'s second conjunct as "attitudinal uniqueness." It is less than obvious to me that there is much work done by this distinction in this specific case; we will see a use for the distinction in a later section, however.

18. Even after clarifying in his (2014) that (UTW) is to be read inter-personally, he persists in using examples geared towards showing that a single person cannot rationally believe permissively.

still much to learn by considering his arguments, however (and I can only assure the reader that the spirit of White's objections will play an important role later in the thesis). Like Feldman, he gives arguments focused on two areas: cases where the same body of evidence supposedly pushes in contradictory directions simultaneously, and cases where epistemic standards (or starting points) allows the agents to draw contradictory conclusions. Central to White's argument, and those of others in the literature, is a worry that the arbitrariness he detects in permissive cases would open the door to rational "belief-toggling". The core idea is this: if a case is permissive, in the sense that the agent can believe P or not- P , and the agent knows this, there seems to be nothing wrong with the agent, while in one of those belief-states, inducing herself to "toggle" to the opposite state. This could happen directly, by taking a pill which straightforwardly induces belief contrary to the one the agent currently holds; or it could happen indirectly, by taking a pill which alters the agents epistemic standards such that she will evaluate the evidence as pointing in the opposite direction from what she currently takes it to.¹⁹ If I've not yet taken a doxastic attitude towards P , I can choose either pill randomly. As a result of this, rational belief does no better than flipping a coin in learning true things.²⁰

Of the first kind, he presents a case similar enough to Feldman's Lefty and Righty case:

How could evaluation of the evidence render me rational in believing that Smith is guilty, if it is not reasonable to believe this already? Surely only if the total evidence supports Smith's guilt. And likewise only evidence supporting his innocence could make it rational to believe that he is innocent. [...]

But the evidence cannot support both Smith's innocence and his guilt. Whatever is evidence for P is evidence for the falsity of not- P and hence is evidence against not- P . Of course, certain elements of or aspects of the total body of evidence might suggest that Smith is guilty, while others suggest the opposite. But it is

19. Talk of such pills, which we might assume produce some very local alteration to the agent's physical brain states, is of course mere thought experiment. We might just as well talk about flipping a coin to choose what to believe or how to handle evidence, but pill-talk respects more closely the intuition that belief is not a voluntary affair; we can't will ourselves into belief, but presumably local changes to agents' brains can cause them.

20. We will deal with the worry over the truth-conduciveness of belief in the next chapter.

incoherent to suppose that a whole body of evidence could count both for and against a hypothesis. So then it is impossible that my examination of the evidence makes it rational for me to believe that Smith is guilty but also rational to believe instead that he is innocent. (White 2005, 447)

Thus, he claims, permissivism is false. If it were true, we would have to allow that the agent is rationally permitted to belief-toggle between believing that Smith is guilty and believing that Smith is innocent. After all, the case is permissive, and either belief is rational, so in toggling I do not come to believe something irrational.

But it is far from clear that any respectable permissivist (or permissivism) would disagree with him. To the contrary, Bayesians, who are on all accounts quite radical about the range of permissible rational doxastic states, would agree with him that a single agent cannot rationally draw contradictory conclusions. So, as indicated, the argument has some force against intrapersonal permissivism, but not the interpersonal variant.

White's "epistemic starting points" cases proceed quite similarly, and are unfortunately also (with one exception we will discuss a bit later) geared towards undermining an intra-personal permissivism very few actually seem to defend. One case of this kind goes as follows:

I follow standards S because I was inculcated with them at MIT. But had I attended Berkeley, I would have been inculcated with standards S' instead. Given my total evidence as input, S and S' deliver the conclusions P , and not- P respectively. ... Now I can imagine myself in a counterfactual situation before graduate school where my sole motive for study is to answer the question whether P . I have all the available relevant evidence, I'm just not sure yet what to make of it. Now I learn that if I attend MIT I will inevitably inherit standards S from my mentors, which given the evidence will lead me to believe P . Attending Berkeley will result in my adhering to standards S' and hence arriving at the conclusion not- P . Now surely the prospect of several years of graduate school will seem rather pointless no matter how passionately curious I am as to whether P . Indeed my prospects for answering whether P to my satisfaction should seem dim. I might as well choose a grad school to attend and hence opinions to hold by a preference for Massachusetts weather, or by flipping a coin. Once

I have filled out the enrollment form for MIT say, I will know that unless something gets in the way, in a few years I will be of the opinion that P . If I am a permissivist, I should take it that I will soon rationally believe P . If this is so, why shouldn't I just believe it now and save myself the time and trouble? But of course it would be absurd to form an opinion on the matter by an arbitrary choice when I don't even know what to make of the evidence. If this is the sorry state I find myself in with respect to answering whether P before I begin my inquiry, then I should judge myself no better off having arrived at a conclusion, if I judge that my adoption of epistemic standards was such an arbitrary matter. (White 2005, 452)

White regards the example somewhat “far-fetched”, but as Titelbaum and Kopec point out: “G.A. Cohen tells a very similar story about how philosophers of his generation formed their beliefs concerning the analytic/synthetic distinction depending on whether they attended Harvard or Oxford” (ms. 14).²¹

The case bears quite directly on the question of epistemic standards raised by Feldman earlier. One might, indeed, imagine that what one is doing in choosing where to go to graduate school is precisely picking one's priors as regards P , or picking standards of abductive reasoning. The problem, then, would be seeming arbitrariness in this choice. Here, again, the permissivist would have to allow for belief-toggling via induced changes to one's epistemic standards. That would at least be the reading of the case most amenable to the defender of Uniqueness. Bayesians, certainly, do not think of themselves as having “chosen” their priors. We do no doubt start out our epistemic lives with something resembling priors, but these are not chosen, and they are almost certainly not a full-fledged probability function (though we should, insofar as we value being rational, aim to approximate one).

We will discuss the legitimacy of alternative epistemic standards at greater length in the next chapter. There is another feature of the case worth focusing on for now, and question whether this would be present in typical would-be permissive cases. For the case, as with all White's cases, is an instance of *acknowledged* would-be permissivism, that is, a case where the

21. We might similarly make the case about a logician choosing between attending graduate school in Australia or Britain, believing that the choice would fix her future belief about the correct logic being non-classical or classical, respectively.

agent believes the evidence is permissive, and that she is free to either believe or disbelieve the same proposition. It is far from obvious that permissive cases would be acknowledged as such by agents – indeed, one possible idea might be that acknowledging a case as permissive constitutes evidence for a moderate position (middling credence or suspension of categorical belief). We might imagine an expansion of White's example, where two agents are placed into the choice of graduate school.²² Which school they choose will, without their knowing it, determine whether they believe or disbelieve some proposition P (by instilling some epistemic standards). A third agent sees and knows all this. She has evidence the former two do not, namely that if someone goes to MIT (Berkeley) they will inevitably inherit standards S (S'), which given the evidence will lead them to believe P (not- P). White's suggestion seems to be that if the latter agent should suspend judgement, the former two should; that the three agents are in the same position epistemically vis-a-vis P . But it is not clear that they are. It seems a perfectly sensible position to hold that the former might, whenever they do end up with some belief as to whether P , be rational,²³ while also holding that the latter agent, who acknowledges the case as permissive, would not be rational in taking a stance of whether P . It is not even clear that the difference between the two would have anything to do with permissivism, because they do not in fact have the same body of relevant evidence; the latter agent's body of evidence includes facts about epistemically irrelevant causal belief-forming factors.²⁴

There are some lingering questions, then, what force any example using acknowledged permissive cases can have in the debate. To my mind, they seem less than obviously relevant. More generally, I suspect that many supposed acknowledged cases are not actually what they seem. A recent paper by Julia Smith (2020) proposes an error theory of sorts for the appearance of acknowledged permissive cases; she argues that what appears to be an acknowledgement that one's peer's *belief* is rational, might rather be an acknowledgement that the *peer* is (or behaves) generally rational(ly), or merely that they are a person worthy of respect or a worthwhile conversation partner. There's also a case to be made that acknowledged permissive cases are

22. I am here paraphrasing Titelbaum and Kopec, ms. 14.

23. As long as the instilled standards have whatever features are necessary for rationality. It is dubious whether a member of a cult which inculcated their members to believe all and only contradictions would have an epistemic standard striving for rationality.

24. It might be, after all, that the latter agent is allowed to pass judgement about whether P , but she would be doing it on the basis of different evidence.

rationally self-undermining: Greco and Hedden (2016) argue that the nature of the role of attributions of rationality is such that attributing rationality to a belief that P is inconsistent with holding the belief that not- P ; it is, in a sense, signaling willingness to defer on the question of whether P . They take this to undermine permissivism generally, but this, as indicated by Smith, seems to me a bit hasty.

White's only example which seems to have any bearing directly against interpersonal permissivism is unfortunately also an acknowledged case, but it is interesting in that it is explicitly targeted against the sort of Bayesianism we will be concerned to defend later. It regards how Bayesian agents update their credence on receiving new evidence. White writes,

Suppose that $[\alpha]$ and $[\beta]$ share . . . total evidence E . $[\alpha]$'s subjective probability for P is x , and $[\beta]$'s is lower at y . [They] each now obtain additional evidence E' , which supports P . $[\alpha]$'s confidence in P rises to x' and $[\beta]$'s to y' , which happens to be equal to the x that $[\alpha]$ held prior to obtaining E' . $[\alpha]$ and $[\beta]$ have each updated [their] convictions appropriately in response to the new evidence. (White 2005, 454)²⁵

Let us assume, says White, as the permissivist is wont to do, that this situation could arise, with both agents being rational in their credences both before and after receiving the latest bit of evidence. Why should α change her credence in P to the higher x' if she views $y' (= x)$ as an equally rational credence in that proposition? The charge, in essence, is that permissivism would (or should) allow agents to functionally ignore evidence: in keeping her credence in P at x (rational because $x = y'$, which is in turn rational) α is, relative to her own priors, believing as if she had never learned E' . I claim that the case has bearing on interpersonal uniqueness, despite its focus on a single agent's rationally permissible doxastic attitudes. The upshot of the case, to me, seems to be that such cases cannot arise because there is only one permissible set of prior credences, which is obviously a version of credal interpersonal uniqueness. If you were a (fixed threshold) Lockean about the relation between full and partial belief, this would in turn entail interpersonal uniqueness about categorical doxastic attitudes.

25. α and β is substituted for "I" and "you", and their possessives, in the original text; the latter makes for clumsy reference in my own comments on the quote ("White" and "the interlocutor"?).

We will return to this problem in the next chapter, when discussing Titelbaum and Kopec's contribution to the debate.²⁶ We have throughout the present chapter raised a number of potential problems, chiefly raising the spectre of deep-seated epistemic arbitrariness, for permissivism. Some have been defused, if only in letter rather than spirit, while others still loom. Before we move on to the next chapter, where we will discuss some positive proposals made in the literature by permissivists, we will briefly look at some more foundational questions regarding the notion of evidence at play in the debate; we will lastly raise, parenthetically, an interesting pro-uniqueness theory by Jonathan Weisberg, before justifying our ignoring it for the rest of the thesis.

1.3 Evidence and evidential support

1.3.1 What is evidence?

There has been, and will yet be, much talk about evidence in the present thesis. A few words should be said about what precisely evidence *is*, both in the context of the debate over permissivism and elsewhere. As a conceptual matter, an agent's evidence must at least be some set of some sort of objects which stand in some sort of supporting relation to other objects. Being rather more straightforward about the matter, I suggest that we think of evidence as a set of propositions known by an agent; that is, *an agent's* evidence is the set of propositions known *by her*. In this I am essentially following Williamson's (2000, esp. chap. 9) thesis that $E=K$, or "evidence = knowledge". There are two central parts to Williamson's thesis, that evidence is propositional and that evidence is extensionally equivalent with knowledge.²⁷ I'll not expend much effort arguing for these, noting merely that the sorts of relations evidence stand in requires it to be propositional, and that it seems to me eminently plausible that it is a prime function of knowledge that what is known is able to stand in the evidential support relation.

It would be interesting to consider how different conceptions of evidence

26. Anyone with any experience with formal (or Bayesian, more specifically) epistemology might feel like this is a non-problem, or at least that the solution is obvious. I can only assure the reader we will get something fruitful out of it in that later section.

27. Note: not conceptually equivalent.

might influence the problem of uniqueness. A lot rests on a sensible reading of two agents having “the same relevant evidence,” and different conceptions of evidence might make this harder or easier to make sense of. My immediate intuitions about the question would be that the finer the grain of what counts as evidence, the less sense there is in speaking of agents having the same evidence, which in turn would trivialize the problem. If an agent’s evidence at some time is simply their entire mental state at that time, for instance, no two agents (absent, perhaps, freak cloning accidents of the Swampman kind) would ever have the same body of evidence. Feldman, alongside Earl Conee, has defended the thesis that justified belief supervenes on bodies of evidence, where evidence is understood along these lines, where just about everything internal to an agent, both occurrent and dispositional, counts as evidence, including her feelings and experiences.²⁸ Evidence, on this account, is not necessarily propositional, and it is hard to see how two agents might come to have the same total body of relevant evidence. As Elizabeth Jackson has put it: “The easier it is for one’s evidence to change, the less interesting the uniqueness thesis becomes; at some point, if we keep expanding our notion of what counts as evidence, uniqueness becomes trivially true” (2021, 319). $E=K$ seems to me a reasonably course-grained substantive theory of evidence, on which talk of the same relevant evidence makes plenty sense; whatever else is going on in two agents’ minds, we can clearly conceptualize there being some relevant set of propositions which they both know.

In the previous section I claimed that acknowledged permissive cases involve a different set of evidence than do unacknowledged but otherwise identical cases. I think this is generally correct, but there’s a remaining problem about what sort of bearing the evidence in question has – or rather, *what* the said evidence has bearing on. It does not straightforwardly have bearing on the truth of the relevant proposition. Rather, it seems to me, the evidence has bearing on the rationality of the belief in that proposition. It is difficult to evaluate how this sort of evidence should impact our further beliefs: after all, most of our beliefs arise from a chain of causal factors some of which will obviously be of a non-evidential nature. It is not clear that evidence of such evidentially irrelevant factors, or counterfactual evidence of

28. The papers in their joint publication *Evidentialism* (2004) develop and defend this view. The relevant reference is to the statement in the introduction that “our version of evidentialism allows that one’s evidence includes one’s feelings and experiences” (2). This is of course not an entry into the debate over permissivism, but is perhaps worth considering as a background to Feldman’s contributions to the latter.

how we would believe differently were these factors different, should lead us to, for instance, moderate our belief by suspending judgement or adopting a middling credence. However, this latter result is not necessary for my case: it suffices that we recognize such evidence as, precisely, evidence, that subjects in otherwise identical acknowledged and unacknowledged cases are in possession of different relevant evidence, and that whatever difference exists between their rational *oughts* or *cans* is not a question of permissivism.

1.3.2 Two-place or three-place?

Apart from the question of what evidence is, there's also the question of what evidence says – or rather, *how* it says what it says. We can formulate this question as being about the notion of evidential support: what does it mean to say that some body of evidence supports some proposition? On the one hand, it is tempting to see the categorical (all or nothing) version of “is evidence for” as a two-place relation; some set of propositions stands in the “is evidence for” relation to some other proposition just in case the former is evidence for the latter.²⁹ That is, where F is a set of propositions:

$$\text{“is evidence for”} \subseteq \text{Pow}(F) \times F \quad (1.1)$$

The “is evidence for” relation works, on this picture, much like a logical consequence relation.³⁰ This is a natural thought: the evidence is what it is and says what it says. It either supports some proposition or it does not. Propositional Uniqueness amounts to a principle that all evidential support is to be understood in this way, as a two-place relation between propositions.

On the other hand, though, it has been argued that “is evidence for” is not two-place, but rather three-place. Only relative to some sort of evidential standard does it make sense to talk about some set of propositions being evidence for some other proposition. Igor Douven makes this case strongly, in his (2009) paper on uniqueness, by considering our best theories of evidence in the philosophy of science. Formally, put this as:

$$\text{“is evidence for”} \subseteq \text{Pow}(F) \times F \times S \quad (1.2)$$

29. Or, for all purposes equivalently, the characteristic function of that two-place relation, taking an ordered pair of a set of proposition and some other proposition to 1 (or “true”) if the former is evidence for the latter, and to 0 (or “false”) otherwise.

30. And in deductive (classical) logical argument, it is indeed the correct relation for evidential support; a set of premises either does or does not deductively support (or entail) a conclusion.

, with S being a set of possible standards for evaluating evidence. Some set of propositions stands in the “is evidence for” relation to some other proposition relative to some evidential standard just in case the first is evidence for the second according to the third.³¹ This would be the most reasonable way of understanding what goes on in a case of abductive reasoning: it is only relative to some way of weighing various theoretical virtues we can say that some hypothesis is the best explanation for some observed body of evidence. Something analogous also goes on in the case of Bayesianism, where the degree of support some set of evidence lends to a hypothesis is determined only relative to one’s priors.³²

This distinction between two-place and three-place evidential support is another way of understanding the distinction we made earlier between objective and subjective Bayesianism. The objective Bayesian holds that there is only one rational prior, so while evidential support is, on this account, *technically* relative to the agent’s priors, there is only one such relevant prior. So, quoting Titelbaum and Kopec, “while evidential support is relative to [a] prior, we need not treat it as an additional input to the evidential support function, since it will always have a constant value (so to speak)” (Titelbaum and Kopec 2019, 209). The subjective Bayesian, on the other hand, will hold that a greater number of priors are rational, with some holding that *any* set of credence which obey the probability axioms is admissible as a prior. In that case, we obviously need a third *relatum* to know whether (to what degree) a set of evidence supports a hypothesis.

1.3.3 Weisberg: Univocity, not uniqueness

Jonathan Weisberg’s (2020b) paper “Could’ve thought otherwise” occupies a fascinating place in the literature on permissivism. Weisberg argues that evidential support is a two-place relation, and that Bayesians and other typical

31. The analogous functional version is obvious; “is evidence for” : $Pow(F) \times F \times S \rightarrow \{1, 0\}$, such that “is evidence for”(X, Y, Z) = 1 iff X is evidence for Y on standard Z, 0 otherwise.

32. Though (unless we understand “evidential support” as simple conditional probability raising, which is indeed categorical) the formalism obviously doesn’t quite fit this case, as we’re dealing with degrees of support; we’d need a four-place relation with a place for degree of support and S as the set of all possible priors on F , or something like the functional version with the range being the unit interval $[0, 1]$. The standard Bayesian formalism, which we will introduce in chapter 2, makes better sense of this.

permissivists err in making evidential support relative to any sort of standards; yet he defends not only interpersonal, but also *intrapersonal*, permissivism about rational belief. He draws on empirical work on the psychology of evidence-processing, finding that all live models in that literature have in common that they are stochastic. Evidence-processing, even in the best of cases, is partly randomized, plausibly to such a degree that a given body of evidence could, depending on how the processing goes, give rise either of two contradictories. Weisberg, then, defends (what we, following Titelbaum and Kopec, called) propositional uniqueness while denying attitudinal uniqueness: a body of evidence uniquely supports or fails to support any given proposition, but a body of evidence does not uniquely determine for a given proposition whether we should or should not believe it.

It is somewhat difficult to place the suggestion in the dialectic. It is, on the one hand, a welcome empirical intervention. On the other hand, however, it is not clear how exactly Weisberg's theory have any bearing on what is, ultimately, a debate over rationality. It might simply be that the evidence-processing module in human brains are fundamentally irrational. Weisberg, indeed, acknowledges that the upshot of his argument might be that uniqueness holds for rationality vis-a-vis evidence, even if it does not hold for human belief formation.³³ We might make a comparison on this point to the 1970's studies by Kahneman and Tverski, and the following body of research, on heuristics and biases. Human beings are remarkably (and systematically!) bad at, among other things, reasoning probabilistically. It would be rather untoward, however, for a theory of correct reasoning, or rationality, to make allowances for this, e.g. by permitting that we judge the probability of a conjunction to be higher than the individual probabilities of the conjuncts.³⁴

Weisberg, then, is to my mind most usefully put in the pro-Uniqueness camp. On this front his arguments against permissive epistemologies do not contribute much (though I will return to them in discussing Schoenfield in the next chapter). The focus is on the alleged arbitrariness and instability he, like White and Feldman, sees as arising in (alleged) permissive cases; and he, again like White and Feldman, focuses on *acknowledged* cases. While empirical psychological work on evidence-processing is no doubt interesting, it

33. This is one of the points of §7.2 of his paper.

34. This latter phenomenon is observed in the case of Linda, who most subjects judged likelier to be a feminist and a bank-teller than just a bank-teller.

plainly has more bearing on psychology (and perhaps behavioural economics) than on the philosophy of theoretical rationality.

Chapter 2

The permissivists strike back

We saw in the previous chapter a worry that allowing evidentially irrelevant factors to influence what we are rationally permitted to believe may land us deep arbitrariness. One response to this, White and Feldman's, is to say that we are often, whenever we follow those irrelevant factors, irrational. Another response, one we will examine in this chapter, is to save our everyday rationality by arguing that these evidentially irrelevant factors are not so damaging to rational belief as White argues, justifying, along the way, some form of permissivism. This latter case has been put most forcefully, and expressed in its strongest form, by Miriam Schoenfield, starting with her "Permission to Believe." I begin the chapter by looking at her proposal. Later, I discuss a somewhat more modest approach by Michael Titelbaum and Matthew Kopec, primarily from their "When Rational Reasoners Reason Differently" (but occasionally from their unpublished "Plausible Permissivism" where the contents of the latter go beyond that of the former). I close out the chapter by reflecting on where we've been and where we're going, situating the discussions in the later chapters within the material in the first two.

2.1 Schoenfield's standards

Schoenfield's argument for permissivism focuses on a supposed range of rationally permissible epistemic standards for evaluating evidence, with the conclusion that different standards sometimes give different outcomes relative to the same body of evidence. We should start, then, by introducing Schoenfield's notion of epistemic standards, and the role they play in the

argument.

Schoenfield's idea of an epistemic standard is best summed up as a set of rules of the form "given evidence E , believe p ," or, equivalently, as a function from bodies of evidence to doxastic states (where a doxastic state can, for present purposes, be identified with a belief set B) (Schoenfield 2014, 199). These rule are, or should be, considered somehow truth-conducive to the agent who holds them. There are, says Schoenfield, 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" (Ibid.).³⁵

On this basis, Schoenfield admirably fends off the objections from White: given some agent α , some set of epistemic standards S , and some body of evidence E , it is not the case that the agent can choose whether to believe or disbelieve some proposition P . E relative to S points α quite determinately in one direction with regards to P , even if a different set of epistemic standards would have pointed her in a different direction. As Schoenfield puts it in vivid metaphor, it is not the case that there is an evidential support dial which points the same agent in different directions at once, but rather several different dials for different agents (or rather, different epistemic standards), each pointing in a single direction (200). There cannot, on Schoenfield's account, be any single set of epistemic standards such that a body of evidence can point an agent to both believe and disbelieve a single proposition; such a set of epistemic standards would obviously not be considered truth-conducive by the agent.³⁶

There remains a larger question about arbitrariness, however. In acknowledged permissive cases, there does seem to be a worry about how an agent can justify her belief in P while recognizing that believing $\neg P$ would be perfectly rational. If either would be rational, why couldn't the agent choose what to believe by taking a randomly selected belief-inducing pill? We encountered this worry under the terminology of "belief-toggling" in our discussion of White. Schoenfield answers by distinguishing, as we did, two different sorts of cases of this nature, one where the pill directly induces

35. Schoenfield's (2019) argues that the value of rationality consists in precisely giving us epistemic standards which seem internally truth-conducive (though that paper employs the terminology of "endorsing cognitive properties" which the agent considers "accurate"), and that it is unclear what reason we might have for adopting any concept of rationality which goes beyond this.

36. Though there might be some interesting questions about e.g. dialethism here.

(dis)belief in the proposition, and one where it induces (dis)belief in the proposition *via* inducing a change in the agent's evidential standards. The first case is the simplest one to answer: taking a randomly selected pill which induces belief in a proposition may cause the agent to believe contrary to her epistemic standards. Indeed, in choosing whether to look at and follow the evidence or take the pill, it should seem to her that taking the pill would carry a quite substantial risk that she ends up believing contrary to her epistemic standards; whereas choosing to look at the evidence would of course force her to follow her epistemic standards.

The second case, where the pill induces belief *via* causing a change in epistemic standards, is more difficult. The question, really, comes down to whether and how, from the vantage point of one set of epistemic standards, we can evaluate other sets of epistemic standards. According to Schoenfield, the agent in an acknowledged permissive case would decline a belief-toggling pill, because

[a]lthough 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. (Schoenfield 2014, 201)

This follows from the idea of epistemic standards as the agent's considered judgement about truth-conducive reasoning. The defender of uniqueness can, however, demand of the agent that she justify her current epistemic standards, in the way we saw Feldman do in the previous chapter: *why* is your chosen set of epistemic standards more reasonable than any other? Can the agent appeal to her own epistemic standards in answering this?

According to Schoenfield, there is no way of independently evaluating epistemic standards prior to those standards themselves. This is no problem, however; at least not a problem unique to the permissivist. Schoenfield turns the charge around: what independent justification does the defender of uniqueness have for her way of evaluating evidence? How could she evaluate her own epistemic standards without referring to those very standards?

Schoenfield concludes:

[T]he problem with this cluster of arguments [concerning arbitrariness] for UNIQUENESS is that they all, in some way, rely 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. (Schoenfield 2014, 202)

As Schoenfield takes herself to have shown, these assumptions are both mistaken. First, the permissivist can justify her belief in a permissive case just like anyone else, by stating that her truth-conducive standard of reasoning determines what to believe; second, there is no special problem for the permissivist in justifying her epistemic standards – absent a knockdown argument in favour of something like a unique Ur-Prior, the proponent of Uniqueness has the exact same problem.

Schoenfield's brand of permissivism is quite strong. Nothing, on her account, rules out acknowledged permissive cases; indeed, she positively embraces them. She takes the truth of permissivism to entail that there is nothing epistemically improper in maintaining one's belief when one learns that it is caused by (some) evidentially irrelevant factors. In particular, there's nothing improper with maintaining belief upon learning that your epistemic standards have been shaped by non-epistemic factors like environment, ideology, or religion, and that different such factors would lead to different standards (which in turn would lead to different beliefs), even if you believe those standards would be rational.³⁷

There is something somewhat worrying about Schoenfield's argumentation on this point.³⁸ On the one hand, she claims there is no perspective from which to judge our standards prior to those standards themselves, and that once in possession of a set of standards we judge that set, and no other, to be truth-conducive, and therefore rational. This is why she does not consider her view vulnerable to the belief-toggling problem. Yet, on the other hand, she seems to assume we can step back from our current standards and assess the role of evidentially irrelevant factors on both ours and others' standards, and judge the rationality of both. Her permissivism seems, indeed, to consist in precisely the idea that there is an epistemic standards-neutral position from which to judge that there are several permitted standards. Once we have performed this "stepping back" and found an alternative set of standards rational, there is no principled reason we couldn't swap our standards. That they are *my* standards does not seem to carry much weight once I've stepped back into a neutral position; as Weisberg puts it:

37. This is the upshot of §§3-4 of her (2014).

38. Something like the following is put forward in (Weisberg 2020b, 7-9).

[T]he whole point of stepping back is to bracket the commitment we're stepping back from, opening it up to revision. When we step back from a mundane belief like P , we do so precisely for the purposes of considering whether to change that belief. [...] [I]f we step back from standard S to find ... that S is rational but so is S^* , we become permitted to adopt either one. (Weisberg 2020b, 9)

The point seems, to me, quite correct. The problem looks to stem from Schoenfield's view being too open to acknowledged permissive cases; this will be something to keep in mind as we approach our formal models in later chapters. Anticipating slightly: we will find that the most developed justification for probabilism about partial belief has the feature of measuring the accuracy of a doxastic state with a *strictly proper scoring rule*, meaning every rational doxastic state finds itself to be the most accurate. It will also be the case that different doxastic states expects itself, but not every other, to maximize expected epistemic utility in its full beliefs (or suspension thereof). This might seem to build into these models a resistance to acknowledged permissive cases.

2.2 Titelbaum and Kopec

Mike Titelbaum and Matthew Kopec's "When Rational Reasoners Reason Differently" serves both as an overview of the debate, the arguments and motivations of each side, and a defence of some versions of permissivism. A lot of the paper rests comfortably within the same mold as my own approach, looking more carefully at exactly what sorts of permissivism are endangered by a given argument, and deflating the danger to their own preferred variant. I will begin by going through a couple of central moments of the former kind, which will have some connection to the formal models we look at in later chapters, before presenting their positive case for permissivism.

2.2.1 Bayesian updating in permissive cases

We left off the section on White with a problem for the permissivists, especially of those of the Bayesian persuasion: why should agent update on their evidence if they find an equally rational agent whose post-evidence credence

in a proposition equals her own pre-evidence credence in that proposition?³⁹ Somewhat starker: a Bayesian agent can be sure that, even if there is no actual interlocutor in such a situation, there exists some probabilistic set of priors which, updated on the evidence, would be equal to her own un-updated credence.

Titelbaum and Kopec (henceforth “T&K” for short) appeal here to the principles of subjective Bayesianism; we’ve yet to discuss precisely what this consists in, and will make use of it later in the paper, so some introductory remarks seem in order. Subjective Bayesian epistemology is concerned chiefly with partial belief, or credences. That belief comes in degrees and differ in strength is a most familiar part of our epistemic lives, and it comes natural to us to represent these differences in degree numerically, or more precisely, probabilistically.

I treat, for simplicity, the credences of an agent as a single probability space. A probability space is a triple $\langle \Omega, F, c \rangle$,⁴⁰ where Ω is a set of possible worlds (that is, ways the world might possibly be), F is an algebra over Ω (that is, a set of subsets of Ω such that: at least $\Omega \in F$; for any $P \in F$, $\Omega \setminus P \in F$; and for any $P, Q \in F$, $P \cup Q \in F$ and $P \cap Q \in F$) and c (for credence, which is what we’ll use probability spaces for) is a probability function over F (that is, a function $c : F \rightarrow [0, 1]$, satisfying the axioms of probability: for any $P \in F$, $0 \leq c(P) \leq 1$; $c(\Omega) = 1$; for any mutually exclusive (or disjoint) $P, Q \in F$, $c(P \vee Q) = c(P) + c(Q)$). Intuitively, F is a set of propositions and logical compounds of propositions, where a proposition should be understood as the set of possible worlds where that proposition is true.⁴¹

39. The situation, somewhat regimented, was this (with E being the agents’ current evidence and E' evidence they obtain): $c_\alpha^t(P|E) = x$, $c_\alpha^t(P|E') = x' > x$; and $c_\beta^t(P|E) = y < x$, $c_\beta^t(P|E') = y' = x$. The question was why α should update to x' rather than stay at $x = y'$, which she recognizes as rational.

40. We might, if necessary, index any of these to an agent or a time. Writing $c_\alpha^t(P)$ to give α ’s credence in P at t is a particularly pertinent construction.

41. The classical logical connectives correspond to set-theoretical operations on sets of possible worlds: $\neg P$ corresponds to the complement, relative to Ω , of the set of worlds where P is the case; $P \wedge Q$ corresponds to the intersection of the set of worlds where P is the case and worlds where Q is the case, which is to say, worlds where both P and Q are the case; $P \vee Q$ corresponds to the union between the set of worlds where P is the case and the worlds where Q is the case, that is, worlds where either P is the case or Q is the case. Entailment, while not strictly speaking a connective, is usefully understood as a subset relation: a proposition P entails another proposition Q just in case the set of

The central piece of rational requirement to the Bayesian (of all stripes) is (Bayesian) Conditionalization. We write $c(P|Q)$ for one's credence in P conditional on learning that Q . We might write

$$\text{If } Q \text{ is the conjunction of all the evidence } \alpha \text{ obtains between time } t \text{ and some later time } t+1: c_\alpha^{t+1}(P) = c_\alpha^t(P|Q) \quad (2.1)$$

to express the requirement that α ought, on learning Q , set her credence in P to equal her earlier conditional credence in P given Q .⁴²

White's challenge, then, amounts to an ask for a justification of this requirement which is compatible with a permissive outlook. T&K channel a typical subjective Bayesian in giving an answer:

A rational agent has a hypothetical prior, and as she receives new evidence she must update accordingly. If her credence in P at the beginning of White's story is x , and that credence commits her to a hypothetical prior which assigns x' to P after the addition of evidence E' , then it would be irrational of her to arbitrarily switch to a different credence at the later time. (Titelbaum and Kopec, ms. 8)

The subjective Bayesian holds that there are several rational hypothetic priors, but that each must update on the evidence according her particular prior. This much is common knowledge to anyone working in the field. There is a further question, however, about what precisely forces me to *stick* to my own priors. This ties into the larger problem of arbitrariness White has pushed on us: if there are several rational priors, why can't I take my cue on P from different ones before and after I receive some piece of evidence?

The proponent of uniqueness would, presumably, attempt to ground Conditionalization as a diachronic rational requirement in the idea that your rational credences at some time given some evidence are the only rational credences at that time given that evidence. That is to say, they might ground

worlds constituting P is a subset of the set of worlds constituting Q ; that is, every P -world is also a Q -world. The requirement that F be an algebra on Ω ensures that F includes the negation of every proposition, and both the conjunction and the disjunction of every pair of propositions.

42. This is a special case of what is often called Jeffrey Conditionalizing, where the evidence obtained does not have to be certain. If α 's evidence Q at time $t+1$ determines a partition $\{q_1, \dots, q_i\}$ of logical space summing to probability 1: $c_\alpha^{t+1}(P) = \sum_{i \in Q} c_\alpha^{t+1}(q_i) * c_\alpha^t(P|q_i)$.

diachronic consistency in Uniqueness. It is not, however, clear that such an approach would provide reason to favour interpersonal rather than mere intrapersonal uniqueness; given *your* priors, there's only one thing *you* are permitted to believe given some evidence.⁴³ Perhaps the Uniquers could try for something stronger: your current rational credences are the only ones permitted given your evidence because they are the only rational credences permitted *anyone* given that evidence. The argument would have to be made before we can judge it, and it is not obvious how one would make it.

2.2.2 Truth-conduciveness and misleading evidence

T&K also take up a more general version of the arbitrariness worry. White takes the goal of belief to be truth, and he takes evidence to be generally truth-conducive. Permissivism, according to him, allows for doxastic behaviour no more reliable as a guide to truth than flipping coins on taking belief-inducive pills.

It is, firstly, debatable (and much debated) whether the goal of belief in the first instance is truth.⁴⁴ Some matters bearing on this question will be presented in later chapters, but we will (following T&K) grant the assumption, in favour of looking at the second point, that evidence is truth-conducive, and that permissivism fails to account for this. Concerns about permissivism vis-a-vis the truth-guiding function of evidence come up at various points in the 2005 article (448-9, 451-2, 455), but the clearest formulations is in White's later (2014):

If there is evidence available strongly supporting one verdict, then it is highly probable that it supports the correct verdict. (315)

In a non-permissive case where the evidence directs us to a particular conclusion, following the evidence is a reliable means of

43. While the case presents itself as defending interpersonal uniqueness, it seems on reflection no more than a restatement of the belief-toggling worry, which only supports intrapersonal uniqueness.

44. Meacham (2013, 1208-14) presents, in discussing White's paper, some ways formal epistemologist have objected to the "truth as the goal of belief" idea, and some ways in which the idea can actually be read as supporting permissivism. Splitting your credence in the result of a fair coin flip 50/50 rather than believing (as it turns out correctly) with full certainty that it lands heads up seems reasonable, but would derive you of the pleasure of believing truly. See also Easwaran (2017) for a discussion of the alternatives to truth as goal in both formal and traditional epistemology.

pursuing the truth. (ibid.)

Common wisdom has it that examining the evidence and forming rational beliefs on the basis of this evidence is a good means, indeed the best means, to forming true beliefs and avoiding error. (322)

The point here might, suggests T&K, be summed up as the claim that “most evidence isn’t misleading” (Titelbaum and Kopec 2019, 215). It is an obvious assumption here that evidential support is a two-place relation: the only way to get an anti-permissive argument from the foregoing is if there is a single fact of the matter about what the evidence say. Only then do we get the worry that permissivism somehow permits agents to believe “contrary to the evidence” (not a quote, but clearly the only way sense can be made of the charge).

For the argument to get off the ground, we will have to grant White evidential uniqueness, at least for a moment. We can see the claim as amounting to the following, then: for most pairs consisting of a set of evidence and some hypothesis, the former stands in the relation of evidential support to the latter only if the latter is true. This would cause some problems for the permissivist, as this suggests, given truth as the goal of belief, there’s only one proper way to respond doxastically to a given body of evidence.

But there’s an alternative, permissivism-friendly, way of taking on board the “most evidence is not misleading” claim. As T&K then understand it, the most defensible read would be something like this:

with respect to everyday, useful hypotheses that come up in the ordinary course of life, most people possess bodies of evidence that generally aren’t misleading. This fact helps explain why we tend to have true beliefs in that domain and are able to navigate the world as successfully as we do. (217)

But none of this, they claim, is objectionable to the permissivist. She can simply take it to point towards the fact that most people have epistemic standards which tend to be reliable in evaluating “everyday, useful hypotheses.” This is a perfectly plausible observation for the permissivist; contrary to White’s worries about permissivism allowing for widespread error or crazy inferences, nothing about that position depends on these being widespread. “Most epistemic standards are generally reliable,” then. This does not pick out a unique standard, as the proponent of uniqueness would have it.

T&K, then, have presented an argument somewhat complementary to Schoenfield's, regarding epistemic standards. They conclude:

If epistemically arbitrary causal factors select standards for you *from a set most of whose members are reliable*, the fact that your standards were arbitrarily selected from that set is no reason to question their reliability. (2019, 218)

The charge that permissivism would work against the idea of truth as the goal of belief is, I think, not just defused, but turned around. While most people get things right most of the time, they also occasionally differ: epistemic standards are misleading to different degrees and with regards to different hypotheses.

2.2.3 The Reasoning Room

As the main part of their positive case for permissivism, T&K present a fairly hefty example they call the "Reasoning Room":

You are standing in a room with nine other people. Over time the group will be given a sequence of hypotheses to evaluate. Each person in the room currently possesses the same total evidence relevant to those hypotheses. But each person has a different method of reasoning about that evidence.

When you are given a hypothesis, you will apply your methods to reason about it in light of your evidence, and your reasoning will suggest either that the evidence supports belief in the hypothesis, or that the evidence supports belief in its negation. Each other person in the room will also engage in reasoning that will yield exactly one of these two results.

This group has a well-established track record, and its judgments always fall in a very particular pattern: For each hypothesis, nine people reach the same conclusion about which belief the evidence supports, while the remaining person concludes the opposite. Moreover, the majority opinion is always accurate, in the sense that whatever belief the majority takes to be supported always turns out to be true.

Despite this precise coordination, it's unpredictable who will be the odd person out for any given hypothesis. The identity of the outlier jumps around the room, so that in the long run each agent is odd-person-out exactly 10 percent of the time. This means that each person in the room takes the evidence to support a belief that turns out to be true 90 percent of the time. (Titelbaum and Kopec 2019, 220)

T& claim, and I'm inclined to agree, that this is a perfectly reasonably set-up, and a plausible example of a permissive case: it is rational for each agent in the case to believe what they do on the basis of the evidence they receive, and that evidence does not determine a unique attitude to the hypotheses raised. Furthermore, it is obviously a case where applying one's method of reasoning to reach a conclusion about some hypothesis is more likely to lead you to the truth about that hypothesis than flipping a coin or taking one of two belief-inducing pills. It is not a conceptual truth *that* this is a permissive case, but intuition seems to point us there, and understanding the case in that light delivers some attractive conclusions.

While the precise nature of the case is somewhat artificial, it is not wildly different from what goes on in, for instance, scientific investigation in the actual world.⁴⁵ The epistemic standards at play are also, by all accounts, truth-conducive in precisely the way White claimed permissive cases cannot be, despite the fact that there are presumably evidentially irrelevant factors at play in causing the agents' beliefs. The agents in the Reasoning Room certainly seem a lot better off than the pill-poppers or coin-flippers White describes. If 90% is not reliable enough, we could up the number of agents arbitrarily high - say the room had a hundred agents with one being wrong on each hypothesis, for a 99% reliability rate (221 fn34).

This is also, interestingly, an acknowledged case of permissivism - to some degree. Each agent has evidence that some other agent, equally informed and rational, has come to a different conclusion on a given hypothesis, but is herself rational in maintaining her belief. However, as T&K point out, if you in *discussing* the hypothesis with one of the other agents find her to disagree, there might be some reason to moderate. They explain this in the following manner:

Given what you know about the distribution of opinions in the room, you should expect before interacting with your colleague

45. T&K mention the groups working on the IPCC report as an example in this regard.

that she will agree with you about the hypothesis. Before interacting you believe H , so you believe eight out of the other nine people in the room also believe H , so you expect a randomly selected peer to agree with you. When you find that she believes $\neg H$ instead, this is a surprising result, which leads you to take much more seriously the possibility that you are the only H -believer in the room. So you suspend judgment on H . (Titelbaum and Kopec, ms. 24-5)

This puts the positive proposal in an interesting place. The sort of acknowledged cases T&K permit are different than what we earlier saw Schoenfield would. They are also less vulnerable to the objections we have seen raised throughout (some of which I've expressed agreement with).

T&K also grant part of the argument from epistemologists of disagreement, like Feldman, that meeting someone who disagrees might put some pressure on you to moderate. I think this a smart move. It is less than perfectly clear how cases like this might arise in less artificial cases than the Reasoning Room, but there is a fruitful and developing literature on how to update on the credences of others (borrowing the title from a paper by Easwaran et al. (2016)), which might contribute.

T&K consider some alternative interpretations of the Reasoning Room example, which might be put forward by the proponent of Uniqueness. The first is simply to deny that the agents in the case are all rational: at most one of them reason from the one true epistemic standard, all the others reach their conclusion by irrational means. It is not clear how the permissivist might respond to this; as noted above, T&K present the example as a way of drawing out some attractive features of the sort of permissivism they favour. "Simply denying permissivism as it applies to the example," they write, "misses the point" (Titelbaum and Kopec 2019, 225).

The second alternative is more interesting. The proponent of Uniqueness might argue that the agents are operating with different total relevant evidence. T&K present, as an example of the sort of evidence which might be posited, "the fact that you have reasoned from the original evidence to H " (226). There will be different facts along these lines for the different agents, hence they are not holding their beliefs on the basis of the same total relevant evidence. While I'm sympathetic to the idea that such facts might constitute evidence, I'm less convinced that they constitute relevant evidence as to whether H . Something along these lines is the response from T&K:

Out of all of an agent's evidence, only what's relevant to a hypothesis may rationally influence her attitudes, and that relation is determined by the agent's epistemic standards. To deny that the Reasoning Room illustrates permissivist commitments, the Uniqueness theorist must establish *in a manner acceptable to permissivists* that all rationally permissible epistemic standards treat facts about one's own reasoning concerning a hypothesis H as evidence relevant to H (Titelbaum and Kopec 2019, 226)

I think T&K get things right here. Their arguments, generally, open up an attractive permissivist position, one on which unacknowledged interpersonal permissive cases seem fairly widespread, and acknowledged cases less so. They make no defence of intrapersonal permissivism, which I suspect is the right call, as far as their own leanings go.⁴⁶

For all this, there's still a quite natural and intuitive suspicion about permissivism. As White puts it in his (2014, 315), permissivism "appears to require a departure from a very natural way of thinking about evidence and rationality." There are still live issues about how to handle disagreement, especially in areas (especially of the sciences) where getting to the truth are most important. A natural, and not unfounded, worry might be that widespread permissivism is undermining the objectivity of those inquires. There is little I can or will do to dispel these worries and suspicions.

2.3 State of play; looking forwards

The first two chapters of the present thesis have examined some key positions and arguments in the literature on permissivism and uniqueness. In chapter 1 we encountered some central worries around the arbitrariness of rational belief on the version of permissivism Feldman and White argue against. In White this took the form of the "belief-toggling" objection: if I in some epistemic situation recognize that there are several rational attitudes I could take to a proposition given my total evidence, there's nothing stopping me from swapping attitudes as I please. The papers from Schoenfield and Titelbaum & Kopec both tackle this objection by, essentially, distinguishing what is

⁴⁶ Though I am intrigued by Douven's (2009, 351) suggestion that sudden flashes of brilliant insight in abductive reasoning might provide something like intrapersonal permissive cases. We will also see that the formal results presented in the next chapter will entail some form of intrapersonal uniqueness; we'll cross that bridge when we get there.

open to the individual agent and what is open to different agents, epistemically speaking. That is, they distinguish between what we call interpersonal and intrapersonal permissivism, and they claim that the arbitrariness worries apply only, at least as presented, to intrapersonal permissivism. The fault in the impermissivist argument, on this defense, seems to be that it runs these together, and claims that if one is true, both are.

Schoenfield and Titelbaum & Kopec make this defense in somewhat different ways, but they converge around the idea that different agents with different *epistemic standards* can legitimately and rationally come to different conclusions on the same proposition given the same evidence. When the notion of an epistemic standard is left floating, perhaps with the minimal description “rules for moving from evidence to beliefs,” the idea seems less than clear; when filled with the more familiar notions of standards of abductive reasoning and Bayesian priors, it seems not only clear, but eminently plausible. I’ve made no attempt at hiding that this is where my sympathies lie in the debate; absent strong arguments in favour of something like Objective Bayesianism

The rest of the thesis will take a somewhat different approach to the question of permissivism. I leave the more general arguments behind, and look rather at how one particular formal model of rational belief might be taken to bear on the problem. I take this to serve two purposes: it is, on the one hand, interesting in itself to see what this independently plausible model of rational belief predicts about the case of permissivism. On the other hand, it is an interesting case study in applying formal methods to more traditional problems, and seeing how to interpret and apply the results.

Chapter 3

The epistemic risk approach

This chapter will take a look at one way of modelling the relationship between rational full and partial belief which focuses on agents' attitude to what we will call "epistemic risk": the respective value/disvalue of believing true/false propositions. These attitudes will affect when a credence is high enough to constitute/warrant/justify outright belief.⁴⁷ It will be shown that if there are different legitimate ways of valuing truth/falsity, there will be different rational distributions of outright belief given the same underlying credences, and given the same evidence. That is, *even given objective Bayesianism*, the *via risk* will deliver interpersonal attitudinal permissivism for full belief.

I begin the chapter by briefly introducing the idea from its original source, William James. I then introduce some of the formal tools we will be employing, borrowing from decision theory. After this, I get to work constructing our model, drawing primarily on Kenny Easwaran's (2016) paper "Dr. Truthlove". After developing the model, we draw out explicitly its permissivist conclusions and discuss Thomas Kelly's (2014) contribution to the

47. The metaphysics of belief, as concerns the question of whether we have both full and partial belief, or whether one is reducible to the other, is a thorny issue. It will be a guiding idea going forward that we have both, though I will not expend much time arguing this. I note merely that we have both philosophical arguments (e.g. in Leitgeb (2017), Buchak (2014), Kaplan (1996) and empirical support from psychology (Weisberg (2020a)) for this idea. If the reader prefers to understand full belief as a simplifying tool in speech, as shorthand for "sufficiently high credence to warrant assertion/action/presupposition," the following will, presumably, be of interest. On the other hand, if the reader thinks credences are mere metaphor, I have very little to offer; this would require a very different kind of argument. (Somewhat ironically, given that the latter reader is the intended audience for Easwaran's paper.)

informal literature on permissivism, which draws some similar conclusions. I then take on some substantive possible worries about the proposal, around the role of evidence, the epistemic legitimacy of the notions employed in the modeling, and the question whether there are in fact several rational attitudes to epistemic risk. I end by briefly discussing Richard Pettigrew's (ms.) risk-based argument for a plurality of permissible prior credence distributions.

3.1 Believing truly, believing falsely

It is epistemological common sense that there is epistemic value in believing true things, and that there is epistemic disvalue to believing things that are false. One particularly vivid description of these two goals and their relation to each other was formulated by William James in his "The will to believe":

There are two ways of looking at our duty in the matter of opinion,—ways entirely different, and yet ways about whose difference the theory of knowledge seems hitherto to have shown very little concern. We must know the truth; and we must avoid error,—these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws. [...]

Believe truth! Shun error!—these, we see, are two materially different laws; and by choosing between them we may end by coloring differently our whole intellectual life. We may regard the chase for truth as paramount, and the avoidance of error as secondary; or we may, on the other hand, treat the avoidance of error as more imperative, and let truth take its chance. [...]

We must remember that these feelings of our duty about either truth or error are in any case only expressions of our passional life. (James 1897, 18, cited in Pettigrew ms., 18)

As James hints at, and as others have later developed further, which of these we consider most important will have great effects with regards to how we form beliefs. Neither can reasonably be taken as all-important: for, were maximizing believed truths the only goal, we could simply believe everything; were the absence of false beliefs the only goal, we could simply abstain from

believing anything. But between these two radical ways of valuing truth and falsehood, there seems a vast range of prima-facie reasonable attitudes.

At this point we will need to introduce some more technical vocabulary and notation. As regards the attitudes towards the value of truth and falsity, we will assign numbers to these values,⁴⁸. We label them R , the value of believing rightly, and W , the disvalue of believing wrongly. We give both as non-zero positive numbers, and give W a negative sign in calculations; while there might be a great range of permissible attitudes to epistemic risk, we think it safe to rule out attitudes which assign a negative value to getting things right or a positive value to getting things wrong, as well as attitudes indifferent to either (for reasons laid out in the previous paragraph).⁴⁹ We include also the epistemic utility of suspending belief on a proposition, which has the value of 0.

We need also some more precise ways of speaking about agents' full beliefs. I will use set-theoretic constructions to model the belief states of agents. The set of propositions believed outright by an agent is labeled B (if necessary indexed to an agent (B_α), to times (B^t), or to both (B_α^t)); we write $P \in B$ to express that the proposition P is in an agent's set of beliefs and $\neg P \in B$ to express that not- P is in the agent's set of belief (that is, she outright disbelieves that P). There will also be propositions P such that $P \notin B \wedge \neg P \notin B$, that is, propositions where the agents suspends judgment as to whether P .⁵⁰ I will at times abuse notation and use B to predicate of an agent that they stand in the belief relation to some proposition; B , then, stands also for

48. With the usual caveats about numerical measurements and the underlying mental goings-on (I discuss this briefly in §4.5).

49. Even with the assignment $\{R = 1, W = 0\}$, where the difference is small, the epistemic benefit of believing a proposition will always outweigh the potential downside to getting it wrong, no matter the probability one assigns to that proposition being true. Anticipating the expected utility machinery slightly: consider $\{c_\alpha(P) = 0.4, c_\alpha(\neg P) = 0.6\}$; the expected value of believing P if true (which is $R_\alpha * c_\alpha(P)$ or $1 * 0.4 = 0.4$) is greater than the expected disvalue of believing it if false ($W_\alpha * c_\alpha(\neg P)$ or $0 * 0.6 = 0$), and is also greater than the expected value of suspending, 0.

50. There is a remaining question of whether there are any propositions such that $P \in B \wedge \neg P \in B$ (or, more realistically, whether $\{P, Q, Q \rightarrow \neg P\} \subseteq B$), that is, whether the agent can believe contradictory or inconsistent propositions. As a baseline, we are concerned with permissivism among rational agents, and it is tempting to think that rational agents do not hold contradictory beliefs. On the other hand, it will be shown that consistency in full beliefs is not a rational requirement for agents on the model presented in this chapter. This springs from pressures from problems like the Preface and Lottery paradoxes.

some subset of the product $A \times T \times F$ of agents, times and propositions. We then write $B(\alpha, t, P)$ to express that P is believed by α at t .⁵¹ Putting this together with the formalization of partial belief in the previous chapter, we can characterize an agent's total doxastic state as the set of her full beliefs and the probability space (a triple $\langle \Omega, F, c \rangle$) representing her credences; a central concern in this chapter will be how these hang together.

With this in hand, we introduce some notions from decision theory. The central calculus of decision theory depends on a utility measure and a probability measure. The expected utility of some act A given some set of outcomes O , a probability function c , and utility function u is $\sum_{o \in O} c(o|A) * u(o \wedge A)$. In our case the relevant acts will be believing propositions, the outcomes will be the truth or falsity of those beliefs, and the probabilities will be the agent's credences towards those propositions. The utilities of believing truly and believing falsely will be given by the values R and W , respectively. So we might write

$$P \in B_{\alpha}^t \quad \text{iff} \quad R * c_{\alpha}^t(P) + (-W) * c_{\alpha}^t(\neg P) > r \quad (3.1)$$

to express the requirement that α believe P (at t) if and only if the sum of, on the one hand, the product of R (the utility of believing truly) and the credence in P and, on the other hand, the product of $(-W)$ (the negative utility of believing falsely) and the credence in not- P exceeds some threshold r (this will, in the following, be understood as a statement about the *expected epistemic utility* of believing some proposition).


3.2 As someone who loves truth ...

That done, we can begin introducing our model for relating full and partial belief. I will be drawing on Kenny Easwaran's approach in his 2016 paper "Dr. Truthlove Or: How I Learned to Stop Worrying and Love Bayesian Probabilities." As the title indicates, the paper comes to the problem of relating full and partial beliefs from a specific angle. Easwaran's explicit aim is justify the use of the mathematical tools of formal epistemology to philosophers who do not believe there is such a thing as credences, or do not think they do any interesting work (that is, belief-first or belief-only epistemolo-

51. Translation is straightforward: $P \in B_{\alpha}^t$ iff $B(\alpha, t, P)$.

gists).⁵² The set-up, then, is to start with an agent with some full beliefs and some attitudes about the relative value of believing truly and falsely, and present a way of modelling her credences (at least on the assumption that she is rational). The aim is admirable, but we will put the theory to a different use, namely to understand how a rational agent's credences and full beliefs ought to cohere, given her attitude towards epistemic risk.

We will, for ease of presentation, be dealing with cases where only three possible worlds are epistemically possible for the agent.⁵³ This gives eight propositions to consider.⁵⁴ One perk of using examples with relatively few propositions is that we can illustrate an agent's space of epistemic possibility visually like this:⁵⁵



The nodes in the figure represent propositions, which are labelled by the set of worlds a given proposition is true at. The lines between the nodes represent, going upwards, subset relations between sets of worlds, which in turn makes up entailment relations between propositions. If w_1 is actual, then the proposition true only at that world is true, as is every proposition reachable by following a line upward. We also have an obvious notion of

52. The paper starts with a framing story about Dr. Truthlove, who loves believing truly and hates believing falsely, and doesn't much care for talk of credences and probabilities.


53. Easwaran prefers to talk about "situations," so as to dispel any metaphysical baggage, and to remain open to agents having different beliefs about identical individuals (Superman/Clark Kent or Hesperus/Phosphorus). Titelbaum (2013) likewise argues that we, for reasons of "fineness of grain," should be "working with degrees of belief that natural language sentences are true in contexts" (35). Being clear that our worlds are epistemically possible worlds for an agent leaves this same room open: while it's metaphysically impossible that some true proposition about Hesperus is not true about Phosphorus, it is not epistemically impossible. The reader is welcome to translate "possible world" into "situation" in the following if it suits her better.

54. Generally, as is well known, n possible worlds give 2^n propositions to consider. The set of propositions for a given set of possible worlds is just the powerset of that set of worlds.

55. Recreation of Easwaran's own presentation.

negation: the proposition true only in w_2 is obviously true only when the proposition true in both w_1 and w_3 is false.⁵⁶

With a representation of epistemically possible space like this, we can fill in an agent's beliefs by putting B in a node representing a proposition if the agent believes that proposition, and using $*$ as a placeholder for propositions not believed (I follow Easwaran in this for now, but we will modify the figures a bit later).⁵⁷ We might then have something like this:



This agent believes the proposition true at all worlds, the proposition true at w_1 and w_2 , and the proposition true only at w_2 . We can now start evaluating how rational this agent's beliefs are. As we do not know, in evaluating rationality, which world is actual, we have to consider how well the agent does, epistemically speaking, in each of the possible worlds. If w_1 were actual, she would believe two truths and a falsehood, if w_2 were actual, she would believe three truths and no falsehoods, and if w_3 were actual, she would believe one truth and two falsehoods. Having refrained from giving specific values for R and W , we write this score as $(2R - W, 3R, R - 2W)$.


56. We also have other logical relations: if $\{w_1\}$ is true, $\{w_1, w_2\}$ and $\{w_1, w_3\}$ and the tautology $\{w_1, w_2, w_3\}$ are all true, so we have conjunction; if $\{w_1, w_2\}$ is true, either $\{w_1\}$ or $\{w_2\}$ is true, so we have disjunction.

57. There are countless possible objections or amendments one might make to the simple image here, each of which is worthy of a thesis unto itself. I'm content to note briefly two of them.

1) We will later be assuming a triple set of categorical attitudes: belief, disbelief, and suspension of belief. These will be held to cover all the attitudes an agent might have to every proposition in the relevant set. Surely, though, there are propositions that are somehow on the agent's mind but are not, in the relevant sense, "entertained," by which I mean no more than to be a candidate for any of the categorical attitudes. What to do about these?

2) And we might, in light of work on reasoning according to non-classical (e.g. strong Kleene) logic ask what to do about propositions which are neither true nor false, or both true and false (conditional on the acceptability of such logics, of course). What does it mean to "get things right" about these propositions? Pettigrew makes some attempts at approaching this question in his (2017)

Having a way of representing doxastic states, and a way of scoring them, we can begin looking at how well different doxastic states do compared to each other. Consider the following two doxastic states:



We borrow a bit of further terminology from decision theory, and say that one doxastic state *strongly dominates* another iff they are defined over the same propositions, and the former has a *strictly higher* score for every possible world, that is, it does better no matter which possible world is actual. We also say that one doxastic state *weakly dominates* another iff they are defined over the same propositions, and the former has *at least as high* a score for every possible world, and a *strictly higher* score in some possible world. The left doxastic state above has the score $(3R, 2R - W, 2R - W)$, and the right has $(3R - W, 2R - 2W, 2R - 2W)$. We can see easily that no matter the values for R and W , and no matter which possible world is actual, DS1 will do better than DS2.⁵⁸ For many pairs of doxastic states, whether dominance occurs will be independent of the values of R and W . This is the case for the above states, for DS2 will score W worse than DS1 in every possible world, and given that W is strictly higher than 0, this will always make a difference.⁵⁹ But for other pairs, whether dominance occurs will depend on the specific values of R and W . This will be the case, for instance, for the doxastic state scoring (R, R, R) and the state scoring $(3R - W, 3R - W, 3R - W)$: which dominates depends on the ratio of R to W .

We can generalize this method of scoring in more precise terms.⁶⁰ We

58. Say $R = 3, W = 4$; DS1 will have the score $(9, 2, 2)$, DS2 will have $(5, -2, -2)$. DS1 does better in every world than (that is, strongly dominates) DS2.

59. Some obvious reasons why DS2 is undesirable is might be that it believes both a proposition and its negation ($\{w_1, w_3\}$ and $\{w_2\}$) and it believes an entailing proposition but not what it entails ($\{w_1\}$ but not $\{w_1, w_2\}$). These are easy ways of making one's doxastic state dominated, given some reasonable restrictions on the legitimate values of R and W .

60. This way of presenting things is from Pettigrew (2017).

want a function EU (for epistemic utility) taking a pair $\langle B, w \rangle$, and giving the epistemic utility of a doxastic state at a world. $EU(B, w)$, that is, gives the epistemic utility of B at world w . This is found by summing the individual utilities of the beliefs in B . That is:

$$EU(B, w) = \sum_{X \in F} eu(w(X), b(X)) \quad (3.2)$$

where $w(X)$ gives the truth value of proposition X at world w — t iff $w \in X$ and f iff $w \notin X$ — and $b(X)$ gives the attitudes towards X — B iff $X \in B_\alpha$, D iff $\neg X \in B_\alpha$, S iff neither of the previous two. The local epistemic utilities, eu , are given by a function $\{t, f\} \times \{B, S, D\} \rightarrow [-\infty, \infty]$. This is simply the values for getting things right, wrong, or suspending: $eu(t, B) = eu(f, D) = R$; $eu(f, B) = eu(t, D) = -W$; $eu(f, S) = eu(t, S) = 0$. Dominance can now be put like this: B strongly dominates B^* iff $EU(B, w) > EU(B^*, w)$ for all worlds w . B weakly dominates B^* iff $EU(B, w) \geq EU(B^*, w)$ at all worlds w and $EU(B, w) > EU(B^*, w)$ in some world w . Applying this to the doxastic states DS1 and DS2 would give the exact result we laid out informally, that the epistemic utility of DS1 is greater at all worlds.⁶¹

Dominance has an obvious epistemically normative aspect: if an agent's doxastic state is strongly dominated by another, she will do worse than the dominating state no matter which is the actual world, and this is a good reason to not be in a strongly dominated doxastic state. If an agent's doxastic state is weakly dominated by another, she will do no better, and possible worse, than the dominating state, which is a good reason to not be in a

61. As we've added disbelief and suspension as distinct attitudes, the score would look different ($EU(DS1, w_1)$ is 6R, for instance), but as Easwaran shows in appendix C of the Truthlove paper (842-5), his two-attitude model and a three-attitude model are for all purposes equivalent, provided the values line up as above (belief in truth = disbelief in falsehood, belief in falsehood = disbelief in truth), and belief in P correlates with disbelief in $\neg P$. The latter assumption, a symmetry condition, might be problematic, were it not the case that we are, ultimately, interested in differences among rational agents. As Easwaran puts it:

“it seems plausible that there is a further requirement that one *ought* to have a doxastic state that satisfies the symmetry condition. Thus, for three-attitude doxastic states, we can define coherence in terms of dominance together with the symmetry requirement. Thus, although there will be three-attitude doxastic states that don't correspond to any two-attitude doxastic states, all of the coherent ones will correspond” (Easwaran 2016, 843)

weakly dominated doxastic state. On the basis of this Easwaran introduces an epistemically evaluative notion of coherence: a doxastic state is *strongly coherent* iff there are no alternative doxastic states defined on the same set of proposition which even weakly dominates it. Rational agents, he claims, ought to have a strongly coherent set of beliefs.

We should turn, now, to how an agent's full and partial beliefs should relate to each other. Easwaran shows that if a doxastic state can be represented as maximising expected epistemic utility for some probability function, given some values for R and W (826-9), then it is strongly coherent.⁶² This is the main result of the paper, and what would allow philosophers like Dr. Truthlove, who do not see the value in reasoning about credences, to appeal to the mathematical formalism developed by formal epistemologists over the past decades.

The interesting part, for our purposes, is slightly different, namely that

62. The central move is a pair of theorems about the relation between dominance and expected value:

[(1)] If doxastic state A strongly dominates B, then for any probability function $[c]$, the expected score of A on $[c]$ is strictly greater than the expected score of B on $[c]$.

[(2)] If doxastic state A weakly dominates B, then for any probability function $[c]$ such that $[c(w)] > 0$ for all $[w]$, the expected score of A on $[c]$ is strictly greater than the expected score of B on $[c]$. (827)

These are quite intuitive: if A dominates (strongly or weakly) B, every term in the sum giving the expected score (3.3 below) of A is at least as great as the corresponding terms in the sum giving the expected score of B (by the definition of (strong/weak) dominance). The probabilities of the possible worlds sum to 1. If A strongly dominates B, at least one possible world where A does better than B is assigned non-zero probability, so at least one term in the sum giving the expected score is strictly higher, so the expected score itself is strictly greater.

The only way for dominance to occur without a (strictly) greater expected score would be weak dominance where the strictly higher-scoring world(s) are assigned a credence of 0. If every possible world has non-zero probability, this is clearly ruled out. So:

If we can find a probability function $[c]$ such that A has maximal expected score for $[c]$, then the first theorem tells us that no other doxastic state strongly dominates A, and if $[c]$ doesn't have $[c(w)] = 0$ for any $[world]$, then the second theorem tells us that no other doxastic state even weakly dominates A. (ibid.)

It's a given for our purposes that the possible worlds have non-zero probability, or else they would be epistemically impossible worlds.

given a probability function and a pair of values for R and W, there will be some strongly coherent set of beliefs maximising expected epistemic utility. That is to say, for a given *credence* function, assumed to be probabilistic, and some values R and W, we can find a doxastic state B such that there is no alternative state B* such that B* even weakly dominates B.⁶³ The expected epistemic utility of a doxastic state B is given by

$$\sum EU(B, w) * c(w) \tag{3.3}$$

, with the sum ranging over all worlds w. The question is, for a given probability (/credence) function and some values of R and W, which set of beliefs is maximising epistemic utility? This set is found by taking each individual proposition and figuring out which attitude towards it maximises expected epistemic utility. This seems intuitive, but it is not immediately obvious why. It follows from the “linearity of expectation,” the fact that “the expected value of a sum of ... variables is the sum of the expected values” (Meyer and Rubinfeld 2005).⁶⁴ Having seen that the expected epistemic utility of a doxastic state is the sum of the expected epistemic utility of that state at every world (3.3), and recalling that the epistemic utility of a doxastic state at a world is the sum of the epistemic utility at that world of the beliefs that make it up (3.2), it follows that the expected epistemic utility of a doxastic state is the sum of the expected epistemic utility of the beliefs which make it up.

63. The relation between expected utility and dominance is less than obvious: Easwaran uses expected utility simply to show when dominance occurs, without assuming normative force for the former (827). This has to do with the specific purpose of his paper. I, and others in the literature, consider maximizing expected epistemic utility itself to be the goal.

64. See also Joyce (1998, 587): “The basic law of expectation is an additivity principle that requires a person’s expectation for a quantity to be the sum of her expectations of its summands.” The expected value of two fair dice, for instance, is 7, the sum of the expected value of each, which is 3.5. This is a simple matter of rearranging terms:

$$\begin{aligned} &\text{Let } T = f + g, \text{ which are defined over situations } s \in S, \text{ and recall that the} \\ &\text{general definition of expected value is } \text{Exp}(f) = \sum_{s \in S} (f(s) * \text{Prob}(s)). \\ &\text{Exp}(f + g) = \text{Exp}(T); \\ &= \sum_{s \in S} (T(s) * \text{Prob}(s)); && \text{(by the definition of Exp)} \\ &= \sum_{s \in S} ((f(s) + g(s)) * \text{Prob}(s)); && \text{(by the definition of T)} \\ &= \sum_{s \in S} (f(s) * \text{Prob}(s)) + \sum_{s \in S} (g(s) * \text{Prob}(s)); && \text{(rearranging terms)} \\ &= \text{Exp}(f) + \text{Exp}(g); && \text{(by the definition of Exp)} \end{aligned}$$


(Meyer and Rubinfeld 2005)

Easwaran puts the finding like this (with “score” to be read as “epistemic utility”):

for a given probability function, we can find a doxastic state that maximizes expected score just by figuring out, for each proposition, whether believing it or not believing it has a higher expected score, and then choosing a doxastic state that has the maximizing attitude (for this probability function) to each proposition. (Easwaran 2016, 828)

The expected epistemic utility of believing a proposition P , in turn, is given by $R * c_{\alpha}^t(P) + (-W) * c_{\alpha}^t(\neg P)$; the expected epistemic utility of disbelieving P is $R * c_{\alpha}^t(\neg P) + (-W) * c_{\alpha}^t(P)$; and the expected epistemic utility of suspending is 0.

Let’s look at an example; here’s a probabilistic credence function, shared between agents α and β :



We can now work through, for each of the propositions, whether believing, disbelieving, or suspending judgment on the them maximizes expected epistemic utility, for some pair of values for R and W . Easwaran, and later Pettigrew (ms.), have worked out that for each proposition, which attitude maximises expected epistemic utility can be figured out by looking at the how the credence assigned that proposition holds relative to a threshold of a certain ratio of R to W . In the cases we are concerned with, where $R < W$,

the thresholds look like this:⁶⁵

$$\text{if } c_\alpha^t(P) > \frac{W}{W+R} \text{ then } B(P, \alpha, t) \quad (3.4)$$

$$\text{if } c_\alpha^t(P) = \frac{W}{W+R} \text{ then } B(P, \alpha, t) \text{ or } S(P, \alpha, t) \quad (3.5)$$

$$\text{if } c_\alpha^t(P) < \frac{R}{W+R} \text{ then } D(P, \alpha, t) \quad (3.6)$$

$$\text{if } c_\alpha^t(P) = \frac{R}{W+R} \text{ then } D(P, \alpha, t) \text{ or } S(P, \alpha, t) \quad (3.7)$$

$$\text{if } \frac{W}{W+R} > c_\alpha^t(P) > \frac{R}{W+R} \text{ then } S(P, \alpha, t) \quad (3.8)$$

I leave the thresholds for alternative risk profiles in the footnotes.⁶⁶ The doxastic states rationally required⁶⁷ by $R = 1, W = 2$ and $R = 1, W = 5$

65. A brief explanation of why the first of these thresholds hold; parallel reasoning goes for the rest. Hold fixed the stipulation that R and W are positive integers and that $R < W$. The expected epistemic utility of believing a proposition P is $R * c_\alpha^t(P) + (-W) * c_\alpha^t(\neg P)$. This will, intuitively, be higher than 0 (the EEU of suspending) when the first term is greater than the second, that is, when $R * c_\alpha^t(P) > W * c_\alpha^t(\neg P)$ (it will, for obvious reasons, also be higher than the utility of disbelieving). But this last inequality simplifies to the threshold in (3.4) (shedding the indices):



$$\begin{aligned} R * c(P) &> W * c(\neg P) \\ R * c(P) &> W * (1 - c(P)); && \text{(substituting } 1 - c(P) \text{ for } c(\neg P)) \\ R * c(P) &> W - W * c(P); && \text{(expanding the right side)} \\ R * c(P) + W * c(P) &> W - W * c(P) + W * c(P); && \text{(adding } W * c(P) \text{ on both sides)} \\ R * c(P) + W * c(P) &> W; && \text{(simplifying the right side)} \\ c(P) * (R + W) &> W; && \text{(factoring the left side)} \\ \frac{c(P) * (R + W)}{R + W} &> \frac{W}{R + W}; && \text{(dividing by } (R + W) \text{ on both sides)} \\ c(P) &> \frac{W}{R + W}; && \text{(simplifying the left side).} \end{aligned}$$

So, the EEU of believing P is greater than the EEU of suspending when the agent's credence in P is greater than $W/R + W$.

66. When $R = W$, the agent should believe a proposition she assign credence greater than 0.5, disbelieve a proposition she assigns credence lower than 0.5, and can do whatever she pleases towards a proposition she assign credence 0.5 precisely. The case of $R > W$ is almost the same, except the agent should never suspend – the risk-inclined should always either believe or disbelieve, even at 0.5 credence. Easwaran stipulates that $R \leq W$ precisely for this latter reason; we will discuss this range of attitudes in the next section, but there is reason to believe that the interesting range lies in the $R < W$ section.

67. I say “rationally required” because I believe we have here a strong argument, based purely on considerations of epistemic utility, for a normative version of a threshold view

are represented below, to the left and the right respectively. We modify the earlier representations by adding D in a node if the agent should disbelieve that proposition, and * if the agent should suspend judgement (see note 61 for some background on this change). Our agent should be opinionated over the set of propositions, that is, she should hold a doxastic attitude (suspension is a doxastic attitude) to each of the propositions.



Our agent to the left is moderately risk-inclined. Filling in the thresholds from (3.2)-(3.6), she should believe every proposition she assigns a credence of $2/3$ or higher, disbelieve every proposition she assigns a credence of $1/3$ or lower, and suspend judgement on any proposition she assigns a credence between the former two. *At* a threshold, between e.g. belief and suspension, she can do either. The agent to the right is, clearly, a lot more wary of making mistakes, both in believing falsehoods and disbelieving truths. Her ratios for belief/disbelief/suspension in some p are $c(p) > 5/6$, $c(p) < 1/6$, and $1/6 < c(p) < 5/6$, respectively (the same goes for credences at the thresholds).

Neither of these states are even weakly dominated. The left state is indeed a paradigm case of how to avoid being weakly dominated: decide on one possible world and believe every proposition true in that world – no other state will do at least as well in every world and strictly better in some world. For the right state, adding further beliefs would be “sticking your neck out”:

as a theory of coherence between credences and full beliefs, both of which I take to be actual. Various other reductionist or eliminativist theories will, of course, take different lessons (of descriptive or metaphysical nature): the belief-first epistemologist in Easwaran’s framing story might say that an agent’s credal state (understood as a useful fiction) is some probability function representing her (actual) full beliefs as maximising expected epistemic utility; the credence-first epistemologist, who might hold the threshold view as a simple metaphysical theory, might say that an agent’s belief set (useful fiction or otherwise (at least) supervenient) just is some set which maximises expected epistemic utility for her (actual) credences.

say you add the belief that $\{w_1, w_2\}$ (and correlatively disbelief that $\{w_3\}$), in the case that w_3 is actual, you'd do worse in that world. And the same goes for other beliefs you might add. So the right one is not weakly dominated, and is therefore strongly coherent. We thus have two different belief sets maximizing expected epistemic utility for the same credence function, which is to say, we have permissivism.

3.3 Spelling out the relevance for permissivism

We've shown in sufficient detail, I believe, that given a legitimate idea of epistemic risk and a plurality of attitudes to it, the same probabilistic credence function can give rise to different strongly coherent doxastic states. Slightly more work needs to be done to connect this explicitly with the debate over permissivism. First, we need to explain where the evidence comes in: permissivism is a thesis about the plurality of permissible doxastic attitudes given some body of evidence. But we've so far only talked about relations between various doxastic attitudes. Secondly, for the foregoing to have bearing on the debate over permissivism, we need to justify connecting closely the notions (from Easwaran) of maximising expected epistemic utility and strongly coherent doxastic states with the notion of rational doxastic attitudes at play in the permissivism debate. Lastly, we need to settle which version of permissivism the epistemic risk approach can be seen to support.

3.3.1 Whither evidence?

Evidence has to some degree dropped out of the picture in this chapter. It should rightly be remarked that the foregoing could be taken seriously without a theory of evidence at all: however we come by our credences, they should bear a certain relation to our full doxastic attitudes. But we are of course working within a broadly Bayesian picture: we assume that we come by our credences by updating our priors on our evidence. The epistemic utility framework we've employed can also be used to justify conditional updating on evidence as a requirement of rationality.⁶⁸ I argued in §1.3.1 that we should understand evidence in something like the way Williamson

68. See e.g. Pettigrew (2016, chap. 14) for an epistemic utility argument for a synchronic principle of plan conditionalization, and his (ms. chap. 6) for an epistemic utility argument for a diachronic principle of Bayesian updating.

proposes in his (2000); this is perfectly acceptable for the Bayesianism we are here working with. The evidence our agents ought to update on, then, is the set of propositions they know.

So far, though, nothing here implies that the account of epistemic risk has specific bearing on the permissivism debate: after all, it might seem, different agents could come to have the exact same credences by updating on different evidence. That these two agents might, due to their respective attitudes to epistemic risk, hold different full rational doxastic attitudes, has no bearing on whether permissivism is true. It is primarily when we grant, for the sake of argument, the objective Bayesian the idea that evidence uniquely determines credences, that we get our most forceful conclusions.⁶⁹

Restricting the evidence under consideration to what we earlier termed “relevant evidence” might go some way towards making the results more applicable to agents describable in subjective Bayesian terms. The situation might be something this: the agents do not have *all* the exact same prior credences, but their priors are such that every proposition *with probabilistic relevance* to some hypothesis P (every P' such that $c(P|P') \neq c(P)$; that is to say, every P' which, if learned to be true, would make the agents change their mind in some way about P itself)⁷⁰, in addition to P itself, receive the same credence. Updated on the same body of *relevant* evidence, then, they would have the same credence in that hypothesis P , and different risk profiles might then make a difference in their categorical attitudes towards P . Situations like this might reasonably be thought to occur in the sorts of cases Feldman and White were concerned with; while subjective Bayesianism obviously leaves any probabilistic prior credences about Lefty and Righty open to the detectives, it would certainly not be a miracle if they ended up with identical ones. Of course, again, if their priors differ along this dimension, it is no great surprise if they come to different conclusions.

One take-away of the present chapter, then, is to be taken conditionally: even if evidential support is a two-place relation, and an agent’s evidence uniquely determines her rational credences, her doxastic attitudes to indi-

69. Though I suppose we might also imagine two agents who just so happen to start out with identical priors, and also happen to receive identical evidence. Even were their situation not mandated by objective Bayesianism, the epistemic risk approach would make their case permissive.

70. The fact that there was a full moon on the night of the murder has no bearing on how likely I consider Jones to be guilty ($c(JG|FM) = c(JG)$), but the fact that a bloodied knife was found in his car does ($c(JG|BK) > c(JG)$).

vidual propositions are not so uniquely determined. The material presented is of interest either way, that is, even on a subjective Bayesian approach, but permissivism about rational doxastic attitudes should come as no surprise on such an approach. It is still a welcome assurance that even if agents, for whatever reason, share priors about some restricted set of relevant propositions, they can come away with different categorical rational doxastic attitudes given sameness of evidence.

3.3.2 The attitudes at play

There is a second obvious problem: what licences the inference from permissivism about doxastic states maximising expected epistemic utility to permissivism about rational belief? Some work needs to be done to make the notions in Easwaran's model epistemically legitimate for the debate we've entered into.

The approach laid out in this chapter falls under a larger literature which has come to be known as "epistemic utility theory." Richard Pettigrew presents the general thesis of that literature like this:

according to epistemic utility theory, the epistemic right—that is, what it is rational to believe—depends on the epistemic good—that is, what it is valuable to believe, epistemically speaking. Thus, according to epistemic utility theory, beliefs and credences can have more or less purely epistemic value given different ways the world is. For instance, according to the version of the theory known as accuracy-first epistemology, which combines epistemic utility theory with the axiological thesis of veritism, a belief or credence is more valuable, epistemically speaking, the more accurately it represents the world. Thus, a belief or higher credence in a truth is more valuable than a disbelief or lower credence in it; and a disbelief or lower credence in a falsehood is more valuable than a belief or higher credence in it. (Pettigrew, ms. 7)

The actual epistemic value of our beliefs is obviously not internally available to us, so epistemic utility theory tells us, analogously to classical decision theory, that what is rational to believe is whatever would maximise *expected* epistemic utility. We showed earlier that this is equivalent to avoiding dominance. There is an obvious analogy here to consequentialist theories in ethics:

the right (what it is moral to do) depends on the good (what is morally valuable).⁷¹ The morally valuable is understood as pleasure or well-being or some such – broadly speaking presentable as some manner of utility. As the moral worth of outcomes is less than perfectly clear to us before they occur, we ought to act so as to maximise *expected* goodness.

There are two parts of the version of epistemic utility theory I am working within, each of which might be questioned about its epistemic legitimacy. First, the idea that the what is rational to believe is wholly dependent on what is epistemically valuable, and that epistemic value is constituted by accuracy – believing true things and not believing false things.⁷² Secondly, that the language of utilities and risk appropriately represents these, and that what is rational to believe is subject to risk-adjustment.

The former might face objections from more traditional epistemologists, who have plenty to say about the epistemic virtues and reliable (perhaps understood modally) cognitive functionings. We have not said much about justification. For one, it is not obvious that there’s anything for me to say here: I’m concerned with rational belief, not knowledge or justification. Insofar as rational belief and justified belief diverge (if the set of propositions it is rational for some agent to believe and the set of propositions that agent is justified in believing are disjoint), that’s no worry of mine. On the other hand, though, what I call rational belief is in some sense justified, at least internally. In addition, insofar as we view truth as the ultimate source of epistemic value, it is not clear any of these other approaches will do any better. Indeed, rationality as the absence of dominance seems to me a perfectly reasonable expression of the idea that truth is the only epistemically valuable thing.

The latter is more tricky: why are attitudes to epistemic risk epistemically legitimate? It is natural when hearing the word “risk” inserted into discussions in contemporary epistemology to understand it as a pragmatic rather than purely epistemic point. The literature on “pragmatic encroachment” is, in some sense, about epistemic risk: the value of the truth or falsity of some proposition is such that believing or knowing it demands more. For some, this is straightforwardly put as a variable Lockean thesis: belief in a high-

71. Or, in the Rawlsian terminology for describing “teleological” ethics: “the good is prior to the right.” Wedgwood (2017) takes this wording on board, and prefers to use “epistemic teleology” for what we have called “epistemic utility theory.”

72. On the credal version of epistemic utility theory, this is understood as having high credence in truths and low credence in falsehoods.

stakes proposition requires belief above some thresholds which is higher than that required for a low-stakes proposition, perhaps because belief is a reason for action, and acting on a high-stakes proposition demands more than on a low-stakes proposition. Without wishing to wade further into the debate on pragmatic encroachment, I'll note only that this is a more pragmatized epistemology than what I aim to conduct here.

The epistemic risk I am concerned with here is purely epistemic. It is the relative values of true and false beliefs as such we are concerned with, not their relation to action or any other pragmatic notion. Insofar as we see true belief as valuable and false belief as disvaluable, and consider different levels of caution rational, talk of epistemic risk seems to me epistemically legitimate. It remains conditional whether there are more than one way of valuing truths and falsehoods (more on this in §3.4), but it is not an objection against this idea that it is anything less than perfectly epistemically legitimate.

3.3.3 What precisely is established?

The approach laid out here was first presented in the informal literature on permissivism by Thomas Kelly in his (2014). In the section “A Jamesian argument for Permissivism”, he rehearses the points laid out at the beginning of the chapter: the goals of believing truth and not believing falsehood are in principle separable, and differences in these “cognitive goals” can lead to differences in rational belief between agents, even given shared evidence. The conclusion he draws is rather careful: it is presumably not the case that “anything goes,” epistemically, but

[s]o long as there are at least some possible cases in which it is reasonable for different individuals to give at least somewhat different weights to the goals, then this can affect how much evidence they should hold out for before they take up the relevant belief. (Kelly 2014, 302)

The conclusions we can draw from the material in this chapter appear quite similarly situated.

At a first glance, then, it might seem we've established what we have called interpersonal attitudinal permissivism for categorical belief. No small feat, that, given that we've granted the objective Bayesian her case. But in fact, looking at the way the calculation of expected epistemic utility translates into thresholds for belief/suspension/disbelief, it turns out that

some credences allow a single agent to either believe, disbelieve, or suspend judgement on a proposition. The uniform probability function on the eight-proposition algebra, given a risk profile of $R = 1, W = 2$, for example, allows an agent to either believe or suspend on all the two-world propositions (and likewise, either disbelieve or suspend on the single-world propositions). At any risk profile with $R > W$, a credence of 0.5 allows the agent to either believe or disbelieve.

This is quite strong. We'll look at whether it is too strong, intuitively, in the next section. First, however, I'd like to consider briefly how the proposal intersects with the other issues we discussed in the first two chapters of the present thesis. We will discuss these issues in greater detail in the next chapter. Firstly: what does the epistemic risk approach predict about acknowledged permissive cases? Insofar as an agent's expected epistemic utility is normative for her judgments about the rationality of beliefs, there will obviously be cases where two agents with the same evidence have rationally drawn contradictory conclusions, where each will judge the other's belief irrational, because not maximising expected epistemic utility. Nevertheless, we must grant Feldman and White, if a case is intrapersonally as well as interpersonally permissive, it might occur that two agents differ in their attitudes towards a given proposition, where each will admit that both her own and the other's attitude is rational, because maximising expected epistemic utility.

Secondly, there's the question of belief-toggling. If a case is intrapersonally permissive, and the agent can adopt either of two (or three!) attitudes to a proposition, what's to stop her from toggling between these, "shopping around," as it were, for which suits her to believe? There's also the further question of indirect belief-toggling, namely how to evaluate agents who change their attitudes to epistemic risk and thereby come to find the demands of rational belief to have changed (say, a credence previously at a threshold being now strictly above/below).

I must admit that these problems sit somewhat uncomfortably in my mind, and I suspect the reader might feel similarly. Especially at risk profiles with $R > W$, where agents can either believe or disbelieve propositions they assign 0.5 credence, we might get acknowledged cases or belief toggling with belief/disbelief, and not just belief/suspension or disbelief/suspension. The next section will hopefully alleviate some worries on this count; what worries linger after that will be the subject of the next chapter.

3.4 Are there different legitimate attitudes to epistemic risk?

Now, we stated the intended take-away from this chapter somewhat: even given objective Bayesianism, we get some measure of permissivism in our full attitudes. But there is an obvious second conditionality to what we are warranted in concluding: the foregoing obviously only holds given a plurality of legitimate attitudes to epistemic risk. So we should discuss what problems there might be in postulating such a plurality, and look at ways we might reasonably restrict which attitudes are legitimate. I do the latter first. It should be noted at the outset that on this point I've no mathematical formulas *proving* the rationality of some attitude to risk. However, the precision with which we've presented the ways different risk profiles will judge the expected epistemic utility of various beliefs, and the range of permissibility the different profiles allows for, lets us evaluate the consequences in a fair amount of detail. It allows us, then, consider quite easily which profiles seem intuitively rational (or at least which seem wildly *irrational*).

Firstly, I submit that profiles with $R > W$ are off the table as possible rational attitudes to epistemic risk. The thresholds for belief and disbelief on these profiles are:

$$\text{if } c_{\alpha}^t(P) > \frac{1}{2} \text{ then } B(P, \alpha, t) \quad (3.9)$$

$$\text{if } c_{\alpha}^t(P) < \frac{1}{2} \text{ then } D(P, \alpha, t) \quad (3.10)$$

$$\text{if } c_{\alpha}^t(P) = \frac{1}{2} \text{ then } B(P, \alpha, t) \text{ or } D(P, \alpha, t) \quad (3.11)$$

The actual values of R and W are irrelevant on these profiles, which might suggest that something has gone wrong. It has some intuitively untoward consequences: it is, for example, more valuable to believe each of two contradictories than to suspend judgement on both (when $R > W$, clearly $R - W > 0$). There is, as seen, no place for suspension of judgement at all on these profiles. This seems quite bad, to my mind. The existence of intrapersonal cases where an agent can either believe or disbelieve a proposition, as occurs with credence 0.5, is also less than appealing.

Things are not much better, I think, with the range of risk profiles at

$R = W$. The thresholds for belief are (3.9), (3.10) and

$$\text{if } c_{\alpha}^t(P) = \frac{1}{2} \text{ then } B(P, \alpha, t) \text{ or } S(P, \alpha, t) \text{ or } D(P, \alpha, t) \quad (3.12)$$

While suspension as permissible at 0.5 credence is welcome, it does not seem we are all that much better off. Intrapersonal cases where belief, suspension, and disbelief are all permissible makes, to my mind, rationality far too arbitrary a thing. Another thing to consider here, which goes for both this and the previous range of risk profiles, is that they seem to make belief an almost absurdly lenient thing, with a credence of 0.5 being enough to believe a proposition. The more traditional literature on belief as credence over a certain threshold (whether this is understood as reduction of the former to the latter or a requirement for coherence between the two) have assumed that the threshold is *at least* strictly greater than 0.5, and presumably closer to the 0.8-0.9 range.⁷³ I believe, then, that the range of reasonable attitudes to epistemic risk we should find interesting is entirely in the $R < W$.

While it might technically be correct to refer to profiles of the $R > W$ kind as “risk-inclined,” and profiles of the $R = W$ kind as “risk-neutral,” I would suggest that the risk-inclined-to-risk-averse spectrum, for epistemic purposes, is actually situated in its entirety in the range of attitudes with $R < W$. As noted, the traditional Lockean threshold is in the 0.8-0.9 range, which would correspond to, at the lower end, $R = 1, W = 4$, and at the higher, $R = 1, W = 9$. I would be hard-pressed to describe both $R = 1, W = 2$ and $R = 1, W = 9$ as “risk-averse.” Rather, $R = 1, W = 2$ seems to me an obviously risk-inclined attitude to epistemic risk – which is not to say I find it clearly unreasonable.

If it is as I claim, and the acceptable range of attitudes lays entirely in the $R < W$ range, we should ask which subset of the latter is indeed acceptable. There are some rather extreme points in the $R < W$ range: though we’ve stipulated that R and W are both non-zero positive numbers, we can get arbitrarily close to the point where the value of getting things right is 0 (think $R = 1, W = 9999\dots$), and arbitrarily close to the $R = W$ point (think $W = 99999\dots, R = W - 1$). We could describe these as the closeness of the ratio $\frac{W}{W+R}$ to 1 and 0.5, respectively.

73. (Douven and Rott 2018, 1100) suggest findings from 0.6180, as a lowest conceivable bound, all the way to a highest non-sceptical bound 0.99, with a cluster of arguments for, and basic assumptions of, 0.9 as a reasonable threshold.

The former would bring us into theories on which belief is credence 1 (or something so close as to be indistinguishable). James, famously, chided William K. Clifford for taking up a position like this, where one would rather “go without belief forever than believe a lie.” While there are interesting theories according to which belief is credence 1 (e.g. Clarke (2013)), these typically do not claim this unrestrictedly, in the way the the epistemic risk model would with this sort of profile.⁷⁴ This seems, to me, too risk-averse in matters epistemic – at least intuitively, belief is conceptually distinct from certainty. There’s nothing untoward in asking someone whether they are *certain* about, or if they just believe, some proposition.

The upshot of this is that the range of acceptable attitudes should keep the ratio $\frac{W}{W+R}$ decidedly below 1. How far below? $R = 1, W = 99$ gives a clean 0.99 threshold for belief, one percent short of certainty. Is this unreasonable? I’m not sure. We should keep in mind, however, that wherever the threshold is, an agent will be permitted to suspend on propositions *at* that threshold. The question of where to set the threshold for belief is also a question of how much evidence we can have for a proposition without believing it. The higher the threshold, the less attractive suspension at it seems to me.

On the lower end, having stipulated that W is strictly greater than R , the threshold ends up higher than 0.5. How much more is required? As mentioned earlier, I’ve no immediate reason to reject something like $R = 1, W = 2$, with its threshold of $2/3$ for belief, but anything lower might seem to be pushing it. Might we then conclude that the range of acceptable attitudes to risk are those mandating thresholds between $2/3$ and, perhaps, 0.95; that is, values for R and W between $R = 1, W = 2$ and $R = 1, W = 19$? I’m not sure, but I could certainly see reasonable people adopting full beliefs in a rational manner based on those profiles.⁷⁵

74. Clarke’s claim is that what counts as full belief is dependent on the context of reasoning, because one’s credence in a proposition changes from context to context, depending on the salient space of epistemic possibilities. So, one can have credence 1 (and therefore full belief) in a proposition in one context, but a lower credence in a different context, without obtaining any new evidence. Others might hold that belief is question sensitive, where the question asked determines some relevant partition of logical space, and that belief relative to that question is credence one relative to that partition.

75. One could, though I will not do so here, run an argument against the application of epistemic utility theory by the permissivism, based on the observation that these restrictions might seem somewhat *ad hoc*; as Steinberger (2019) points out, they certainly do not spring out of considerations of pure epistemic utility. The idea, then, would be that we cannot legitimately restrict the range of admissible risk profiles to those where

This is all assuming there is indeed a plurality of legitimate attitudes to risk. What might the counterargument to this be? There might be something about the nature or aim of belief which makes it such that there is only one rational way of valuing truth vis-a-vis falsehood. It is often said that belief aims at truth.⁷⁶ Might it do so in a way which mandates a single unique pair of values for (what we represent as) R and W, which might,⁷⁷ in turn, give us attitudinal uniqueness? We oughtn't dismiss the idea out of hand, but I think the case for a variable Lockean threshold is very plausible. Digging through the entirety of contemporary epistemology for claims about the nature of belief which might translate into a claim about the uniquely rational attitude to epistemic risk is, anyhow, a task outside the scope of the present thesis. I note only as a point in support of my view Wedgwood's analysis of the "belief aims at truth" idea, where he argues that while a belief is "correct" *qua* belief only if the proposition believed is true,⁷⁸

The fundamental epistemic norm of correct belief, as I have formulated it, does not determine any unique way of balancing the value of having a correct belief about p against the disvalue of having an incorrect belief about p . (Wedgwood 2002, 274)

This idea fits well with my approach in the present thesis.⁷⁹

3.5 The *via risk* for credences

In his (ms.), Richard Pettigrew presents an argument for a plurality of reasonable prior credences, based on the idea of epistemic risk we've employed in this chapter. Establishing such a result is somewhat overkill for our argumentative situation, given that we've conditionally granted the objective

$R < W$, and therefore cannot rule out the very bad forms of permissivism (especially intrapersonally).

76. In various, more or less metaphorical senses, from Aristotle to Sosa.

77. "Might," not "would," for some values of R and W permit more than one attitude to some proposition, and who's to say the unique pair of values isn't one of those?

78. That is to say, it is constitutive of belief as a mental state that one ought to be in that mental state only towards true propositions.

79. The discussion in that paper, on pages 274-6, about the range of admissible "precisifications" of the concept of rational belief is essentially a different manner of formulating the present discussion about the range of admissible attitudes to epistemic risk.

Bayesian her case; I nevertheless want to sketch the argument – if only to bolster the *ethos* of the epistemic risk approach.

One central project in contemporary formal epistemology has been to find non-pragmatic arguments for why our set of credences should be probabilistic, that is, why they should follow the axioms of probability theory. One central such argument is that probabilistic credences are more *accurate* than non-probabilistic credences.⁸⁰ In the utility theoretic terminology we’ve developed: non-probabilistic credences are guaranteed to be *dominated* by probabilistic credences; adopting non-probabilistic credences ensures you do worse, that your credences are farther from the truth, no matter which way the world turns out to be.⁸¹ This is a mark of irrationality.

The accuracy (*via* epistemic utility theory) argument proceeds like this:⁸² we define the “vindicated” credence function $v_w(x)$, which assigns credence 1 to x iff x is true at w , and 0 otherwise. We then define a measure of the distance between any other credence function c and v_w . The typical such measure (*squared euclidean distance*) is the following:

$$d(v_w, c) := \sum |v_w(x) - c(x)|^2 \quad (3.13)$$

, where the sum ranges over all propositions x assigned credence by v_w and c . Assuming now that the epistemic utility of a credence function at a world is a function of the proximity between it and the vindicated credence function, we get the following, one variant of the so-called Brier score:⁸³

$$B(c, w) := 1 - d(v_w, c) = 1 - \sum |v_w(x) - c(x)|^2 \quad (3.14)$$

, with, again, the latter sum ranging over propositions x assigned credence by v_w and c . We’ve already looked at the idea of strong/weak dominance as doing worse in all/at least one possible world; this directly applies to our measure. The argument, then, goes: if a credence function c violates

80. This is contrast to the seemingly pragmatic arguments, revolving around Dutch Books and the like, which might not establish more than “if you value money, like gambling, and don’t want to be played by a clever bookie, you should bet according to probabilistic credences.”

81. The *locus classicus* of the approach is Joyce (1998)

82. The presentation is from Pettigrew (2013)

83. The Brier score is sometimes given as a measure of *inaccuracy* instead of the measure of accuracy given below. By those standards, then, the below is $1 - B(c, w)$.

probabilism, there exists some c' such that $B(c, w) < B(c', w)$ for all w and there exists no c'' such that $B(c', w) < B(c'', w)$ for any w .⁸⁴

One interesting feature of credences on this approach appears when we compare different probabilistic credences. On every acceptable way of scoring sets of credences, it turns out that every probabilistic set of credences expects itself to do better than other probabilistic sets of credences. When we examine the *expected epistemic utility* of these credence functions,⁸⁵ we find that each expects itself to score higher than any other. These scoring methods are, in the parlance, *strictly proper*. So, given something like an objective Bayesian evidential probability, obtained from updating the single permissible prior on evidence, expected epistemic utility will be maximised only by matching one's credences to that evidential probability. Thus, Pettigrew:

[I]f we follow Kelly's Jamesian lead and encode our attitudes to epistemic risk in our epistemic utilities, and if only strictly proper scoring rules can be legitimate measures of the epistemic value of a credence, and if we retain the requirement that we choose only options that maximise expected utility from the point of view of the evidential probabilities, we will not obtain permissivism about credences. After all, if r is the evidential probability of X , then assigning credence r to X uniquely maximises expected epistemic utility. It is thus the unique rational response to that evidence. (Pettigrew, ms. 26-7)

The upshot of this is that we must, if we are to obtain permissivism about rational credences (from an epistemic risk approach, that is), encode our attitudes to risk somewhere else.

This somewhere else turns out to be the decision rule for "choosing" priors when we have no evidence to base this choice on.⁸⁶ Pettigrew draws on a literature of risk-adjusted decision rules in decision theory, where different

84. The latter part specifies that the dominating credence function is itself not dominated. See e.g. chapter 2 of Pettigrew (2016) for a discussion of why "undominated dominance" is a better mark of irrationality than dominance *simpliciter*.

85. The expected epistemic utility of a credence function c is just 3.3 adapted for credences, so: $\sum B(c, w) * c(w)$, with the sum ranging over worlds w .

86. The move towards encoding attitudes to epistemic risk in the procedure for generating priors is also motivated by an argument in Horowitz (2017) with the following upshot: epistemic utility theory with strictly proper scoring rules mandates updating by conditionalization, so the only way for agents with the same evidence to have different posterior credences is for them to have started out with different priors. And if they have different

weight is given to different options on top of just the probabilities. In particular, he is interested in a family of decision rules called the Generalized Hurwicz Criterion, which weighs the worse outcomes of an option differently from the better ones. It, in fact, assigns a weight for every outcome of an option: the worst, the second worst, ..., the second best, the best.⁸⁷ In the epistemic case, the options are different prior credences, and their outcomes are the epistemic utility of that set of priors at a world as measured by the Brier score. We then calculate the risk-adjusted utility of the options and choose whichever comes out best. Here's Pettigrew again:

We measure the value of each option by its generalised Hurwicz score, which we calculate as follows: we line up the possible utilities that the option might obtain for us, from best to worst. And then we apply the appropriate weights to each of these utilities—our first generalised Hurwicz weight to the best utility, second weight to the second-best utility and so on. And then we add up the weighted utilities. We then prefer one option to another just in case the generalised Hurwicz score of the first exceeds the generalised Hurwicz score of the second. And we say that an option is irrational if there is some alternative that we strictly prefer. The result is that the rationally permissible options relative to this decision rule are precisely those that maximise this generalised Hurwicz score. ((Pettigrew, ms. 67)

At the very small scale, we can imagine an agent choosing her priors on a partitioning of epistemically possible space into two possible worlds, w_1, w_2 . Let's say she assigns a greater weight to the best-case option than the worst, such that $\lambda_1 = \frac{3}{4}, \lambda_2 = \frac{1}{4}$. That is to say, she is risk-inclined (this is just what it means to assign a greater weight to the best outcome than the worst). There are a number of possible credence functions on these two worlds – indeed, any combination of assignments of unit reals to the two worlds is a possible credence function. Only a subset is of interest, the probabilistic ones. Of these, only two are found to maximise expected utility after risk-adjustment: $c(w_1) = \frac{3}{4}, c(w_2) = \frac{1}{4}$ and $c(w_1) = \frac{1}{4}, c(w_2) = \frac{3}{4}$.⁸⁸

priors in the first place, their attitudes towards epistemic risk has nothing to do with their different posterior credences.

87. It can actually assign the same weight to each option in the preference ranking, in which case the GHC gives the same outcome as objective Bayesianism.

88. The score of the former would be $B(1, \frac{3}{4}) + B(0, \frac{1}{4})$ at w_1 and $B(0, \frac{3}{4}) + B(1, \frac{1}{4})$ at w_1 .

Given a variety of attitudes to risk (which translates to a variety of Hurwicz weights), and assuming the Generalized Hurwicz Criterion is the right decision rule for picking priors, there's a variety of permissible sets of prior credences. Pettgrew defends a quite broad version of interpersonal permissivism: for any probabilistic prior, there will be some GHC weight for which that prior will be maximising risk-adjusted utility. And as we saw with the small-scale example, the approach licenses any permutation of a probabilistic utility-maximising set of priors, so we get a pretty extensive intrapersonal permissivism as well.

Since the former is higher, the utility of this prior is $\frac{3}{4}(B(1, \frac{3}{4}) + B(0, \frac{1}{4})) + \frac{1}{4}(B(0, \frac{3}{4}) + B(1, \frac{1}{4}))$. This is the exact same score as the latter prior, since for that one, w_2 would be the best case and receive higher weight.

Chapter 4

Intrapersonal permissivism, acknowledged permissive cases, and belief toggling

The first two chapters of the present thesis surveyed a battery of arguments against permissivism, and some permissivist answers. The problems revolved around, generally, worries over the arbitrariness of rational belief on the permissivist picture. In White's paper, this took the form of the "belief-toggling" objection: the permissivist is committed to allowing that agents "toggle" their full beliefs in permissive cases – after all, both (or all) the doxastic attitudes are rational, so nothing, it seems, mandates taking and holding one over the other(s). Elsewhere, the problem takes the more general form of a worry over what it means to acknowledge someone else as rational: it is plausible that judging someone's belief as rational entails some sort of deference to that belief – after all, why else would one consider it rational? The epistemic risk approach makes some interesting predictions about these worries; this chapter will, by way of an epilogue to the previous chapter, briefly consider these. This will also be a way of discussing further the more philosophically oriented possible objections to my approach.

4.1 Intrapersonally permissive cases: arbitrary adoption and belief toggling

On the model we've considered, with the suggested restrictions on values of R and W , belief *at* the $\frac{W}{W+R}$ and $\frac{R}{W+R}$ thresholds will licence, respectively, belief or suspension and suspension or disbelief. These cases are, that is, intrapersonally permissive: the evidence⁸⁹ does not determine, for the single agent, which attitude to take. Either will, in both these cases, maximise expected epistemic utility. So, we might, along with White, ask: If you've not yet taken up either, might you choose which to believe by flipping a coin? And once you've chosen one, what is to stop you from simply toggling to the other (say, if that suits you better for practical, political, or ideological reasons)?

The simple, but perhaps unsatisfying, answer is that the model is not in the business of telling anyone which doxastic state they should actually be in. It tells us, rather, which doxastic states we ought not to be in (namely dominated ones) if we are to be described as rational according to that model (and the broader epistemic utility theory underlying it). Much like in traditional decision theory, agents should in principle be indifferent between options which are tied in terms of expected value.

There might be additional reasons, either rational requirements or psychological facts, explaining why an agent, once she has adopted some doxastic attitude to a proposition, shouldn't or wouldn't or couldn't just change that attitude, even if it might be rationally permitted. Having some attitude to some proposition might function as something of a cognitive anchor: the agent might need some overriding reason to step away from her current attitude. That some alternative attitude to that proposition is (also) rational wouldn't be overriding in the required sense.

On a less psychological note, one might wonder whether considerations from the debate over "epistemic conservatism" would be of any help here. The idea of such conservatism is, essentially, this: without an overriding reason *to* change one's belief, staying put and holding on to one's current beliefs is the rational thing to do.⁹⁰ It is far from clear to me that such a

89. "Evidence," when used in this manner in this chapter, is to be understood, continuously with §3.3.1, as credences updated on evidence, on the assumption that we can make sense of sameness of evidence on this picture.

90. As Foley puts it, it "is the view that a proposition acquires a favorable epistemic

position could be directly justified by considerations of epistemic utility. Insofar as our adherence to the latter springs from some version of *veritism* (or “accuracy first”), there’s plainly little reason to care for such a conservative principle. It would, then, have to be a supplement to the theory. Is it a welcome supplement? I am not so sure. It seems to me rather unmotivated from the viewpoint of the epistemic utility theorist; such a supplement would do nothing to promote accurate belief, and might rightly be suspected of being an *ad hoc* solution to a particular challenge.

While there are, then, possible answers of both psychological and more rational kinds to why we do not (and perhaps why we should not) directly toggle our beliefs the way the impermissivist claims would be possible on a permissivist view of rationality, neither is straightforwardly open to the epistemic utility theorist. It is not, however, clear that this is a particularly bad thing;⁹¹ rationality is a purely internal epistemic notion, and if what is internal to the agent underdetermines which attitude is rational, that is simply how it is. There are other dimension on which to evaluate agents’ epistemic doings, like the various other notions of epistemic virtues (broadly speaking, including use of particular cognitive faculties) in the literature. These, presumably, would not allow for intrapersonal permissivism, while still permitting differences between agents.

4.2 Belief-toggling *via* standards

We can apply another consideration from White here as well. He distinguished, recall, two ways in which an agent could belief-toggle. One is the direct way, outlined earlier. But the other way is more indirect, toggling belief *via* changing one’s epistemic standards such that the other permissible attitude is rational for the agent. Might the agent, then, swap up her attitudes to epistemic risk as a way of avoiding taking up some belief she, for whatever reason, does not feel like having?

It has been assumed that the attitudes are representable in a numerically precise fashion, as precise values for R and W, an assumption which might make the worry over where they come from and how they might change more

status for a person simply by being believed by him” (1983, 165).

91. I do think such belief-toggling would be bad, but not obviously for reasons relating to rationality as presently understood. It would, for instance, make practical reasoning, where belief might be understood as “taking as true for some purpose,” a nightmare.

pressing. That assumption is not needed for the model to function as we want it. It suffices that the agent values truth and disvalues falsity, and that the epistemic value of a doxastic state can be evaluated solely by tallying up the truths and falsehoods believed. Easwaran shows how this project might be carried out in an appendix to the Truthlove paper (pp. 838-842). As he puts it there (with ‘attitudes’ understood as the doxastic attitudes), “attitudes and scores are individuated entirely in terms of the contribution they make to the overall normative status of doxastic states” (838). We can, in many cases, “read off” an agent’s risk profile from her the way she would compare the epistemic worth of different doxastic states.⁹² Thus understood, the model makes better sense of the much less precise actual mental goings-on in the agent.

Now, how can we answer the question of whether the model opens the door to belief-toggling via standards? For one, I think the general thrust of Schoenfield’s argument on this front applies well enough. Adopting different attitudes to the value of truth and falsity would lead me to believe in a manner contrary to what I currently think respects the norms of rational belief. We might even pseudo-formalize this claim: by the lights of my current credences and risk profile, the expected epistemic utility of the doxastic state I should adopt after changing attitudes to truth and falsity will be lower than the expected epistemic utility of my current doxastic state. Recall the our example credence function and the illustration of how two agents’ attitude to epistemic risk resulted in doxastic states from §3.2 (p. 56). From α ’s point of view β does not maximize expected epistemic utility, and if the former were to consider changing her standards to the latter’s, she should expect herself to do worse than she currently is. That is why we shouldn’t do belief-toggling *via* standards.

This is at least a *reason not to* change up your epistemic standards, grounded in the normativity of rational belief. It says nothing about how to evaluate actual changes in standards: there will be some period where the agent’s full beliefs are irrational, but presumably she could just reorient her

92. So, for example:

If every doxastic state is ranked equally to any state that gets exactly two more propositions right and one more proposition wrong, then this implicitly defines $W = 2R$. If every doxastic state is ranked equally to any state that gets exactly three more propositions right and two more propositions wrong, then this implicitly defines $W = 3R/2$. (Easwaran 2016, 844)

doxastic state to cohere with her credences, and once again be maximising utility. I don't think this is a bad result: it seems plausible that we might indeed change our attitude to risk over time and in different situations. This might be epistemically legitimate, such as in the case where some agent realizes she's too cautious in believing based on her evidence (someone with a risk profile of $R = 1, W = 99$ should perhaps come to this conclusion).

Perhaps this ought to be treated in a way where it's not truth and falsehood as such which has utility, but that the truth and falsehood of different propositions might carry different weight. Easwaran makes some steps in this direction in appendix G of "Truthlove," as does Kevin Dorst in his (2019).⁹³ If this is to be of interest for the project there (and also here), there needs to be a non-pragmatic reason for valuing the truth/falsity of some propositions differently from others.

4.3 Acknowledged permissive cases: considering others' standards

We've argued that epistemic utilities are epistemically legitimate: rational belief for an agent is whatever maximises expected epistemic utility for her, and agents ought to be rational in this sense. We also argued that agents can rationally come to hold different attitudes to some proposition on the basis of the same evidence, because different attitudes maximise expected epistemic utility for the different agents. An obvious problem here is what happens when these different agents come to learn that they differ in their attitudes to some proposition.

Two sorts of acknowledged permissive cases seem to press themselves on us: firstly, the case where two agents with *the same attitude* to epistemic risk take up different conclusions in an intrapersonally permissive case; secondly, the case where two different agents take up different come to different conclusions because they have *different attitudes* to risk.⁹⁴ Are the agents

93. Dorst takes off from the observation that "it is independently plausible that the value of having a belief (an answer to a question) is a proposition- and context-dependent affair. A given inquiry is driven by a point, a purpose, a question— it is directed toward certain types of answers, and away from others. Thus different contexts of inquiry provide different epistemic priorities" (2019, 188). I must admit to being uncertain how properly epistemic these priorities actually are.

94. Less interesting is the case of A coming to the conclusion licensed by B's epistemic

obligated, in either, to judge the other's belief rational?

The first case is in some sense the most worrying. Insofar as an agent's expected epistemic utility is normative for her judgments about the rationality of beliefs: if someone takes up a doxastic attitudes I judge to maximize expected epistemic utility, I ought to judge that attitude rational. But worries arise here, of the sort we briefly encountered when discussing White, over whether judging an attitude rational shouldn't amount to deferring to that attitude, even if (or especially when) it differs from one's own. This is the argument Greco and Hedden put forth in their (2016) paper on the social role of judgements of rationality, which is precisely to signal a commitment to defer.

As stated, the model says nothing about which doxastic state one should adopt. If two states have the same expected epistemic utility, it says nothing about which you should prefer, nor anything about what to do when encountering someone who from your vantage point maximises EEU but differs in how they do so. We are, again, left to argue from outside the model. I do not think we can get away from admitting that each agent has to judge the other's belief rational; to do so would be to abandon expected epistemic utility as epistemically normative. My intuitions would be, then, that judging rationality is not simply, or not always, a way of signaling deference, at least with deference understood (as Greco and Hedden do) as straightforwardly adopting the belief judged as rational.⁹⁵

The second case is troubling in a different way. I take it as obvious that two agents who take different attitudes to some proposition based on the same evidence because they have different risk profiles do not have to consider each other's *belief* rational. That would amount to denying the claim two paragraphs ago that an agent's expected epistemic utility is normative for her judgments about the rationality of beliefs. But there has to be something to fill in the blank in "I personally believe that *P*, but I see that you are ___ in suspending judgement, given that you are a bit more cautious than I am." Some of the considerations raised about acknowledged cases in §1.2 are

risk profile, despite A's profile recommending the opposite conclusion. Agents come to the right conclusion for the wrong reasons all the time, and there's nothing special with the current approach which requires dealing with this in a special manner.

95. Titelbaum and Kopec (ms. 27fn) argue something similar, that deference in the sense required by epistemically evaluate language, like judging someone's belief to be rational, might mean "incorporating the rational agent's position into one's own deliberation, or moving one's own attitudes somewhat towards the other agent's."

salient here. Julia Smith argues that what appears to be acknowledgments of the rationality of some *belief* in a permissive case might instead be acknowledgments about one's conversational *partner* being generally rational, generally reliable, or just someone worth discussing issues with.

I think something like this is the correct call to make in these sorts of cases. I might judge a person rational if I recognize her to reason consistently by standards which maximise her expected epistemic utility, without thereby granting, in interpersonally permissive cases where we differ, that some individual belief of hers are rational. This seems to me the only way we get any use out of the term "rational" as applied to agents. Otherwise, rational agents might only be those who have the correct set of epistemic standards according to some vantage point. Which vantage point? The proponent of uniqueness would have to say the objective Ur-Prior, or the objective reasonable way of weighing abductive factors, or ... If we are not amenable to uniqueness, it is unclear which precise vantage point we should favour.

The take-away might instead be that only the first sort of acknowledged permissive case licenses the use of "rational." Perhaps the descriptor for an agent one disagrees with but recognizes to reason consistently from some set of epistemic standards is "reasonable," instead.

4.4 Some take-aways and methodological considerations

Some of the material in this chapter is, admittedly, somewhat disappointing. We have been unable to rule out, in principle, belief toggling in intrapersonally permissive cases. We have been unable to dispense with any negative evaluation on the agent who changes her attitudes to epistemic risk and promptly works to restore her total doxastic state to newly utility-maximising status. And we've had to settle for a fairly loose and, as it were, permissive sense of "rational" as applied to agents, to understand the status of acknowledged permissive cases. While these are less than obviously attractive, it does not seem to me that they bring with them any of the unwanted features of permissivism White and Feldman warn about; rational belief does not seem unduly arbitrary.

It should of course also be noted that that we've granted the impermissivist the assumption of objective Bayesianism throughout these two last

chapters; no small assumption, that. A thoroughgoing epistemic utility theory for both credences and full belief would not grant that, as shown in §3.5. The situation described as permissive is not, on this latter comprehensive theory, impossible: surely two agents can have the same priors (because they “chose” them according to the same decision rule) and the same current credences *via* the same process of Bayesian updating, but for whatever reason differ in how they weigh the (dis)value of full belief in truth and falsehood. This would however, intuitively, be a much rarer occurrence than we’ve presented it as.⁹⁶ Interestingly, the intrapersonally permissive cases might be more common, as, in principle, any set of current credences might, for some epistemic risk profile, give rise to such permissivism.

I want to close out the chapter by considering briefly the role of formal models in answering more traditional epistemological problems. As noted throughout the previous chapter, not every part of the problem is solvable with formal methods. We’ve no mathematical result *proving* that rationality is strong coherence, or adherence to probabilism, or ...; we’ve no results *proving* which values of R and W are rational or legitimate; we’ve no results *proving* that epistemic utility, understood as believing truth and not believing falsehood, is the only thing of epistemic value.

What we do have are models of phenomena with various properties, and we then have to do work connecting those properties with philosophically interesting concepts. Firstly, we should ensure that what the models describe actually maps on to the real phenomena we are trying to understand. My goal in this thesis has been to examine a putative principle of rationality, that evidence uniquely determines rational belief. It does not seem to me an objection that we are working at a high level of abstraction, or with too idealized doxastic states.⁹⁷ Secondly, we have to argue that the more or less mathematical or logical properties of the model map on to attractive properties of the real thing, well enough to be norm-giving. Strong coherence

96. But: see again the point about same *relevant* evidence in §3.3.1.

97. There are obvious questions about the idealized nature of formal epistemology more generally: there is not a one-to-one correspondence between the entities we posit in the formalism and the physical brains of real-life epistemic agents; there is no “belief box” corresponding to the set B , and there are no physically realized 0.7-degreed mental events corresponding to what we call a credence with strength 0.7. But in so far as an agent’s belief can, at least, ideally, be “bundled up” in a set, and in so far as the scale of strength of belief observed in everyday epistemic life maps onto something like the unit interval, we are working with the correct tools for the job.

qua absence of even weak dominance is, for instance, a plausible candidate for a condition for rationality because dominance entails the agent recognizing that another doxastic state will do at least as well as, and possibly better than, yours. This seems to me a paradigmatic example of what we want rational justification of belief to do. This comes in tension with a different strong intuition, that a rational belief set should be consistent, which the Easwaran model does not ensure. We then have to set these intuitions against each other.

The model presented in the previous chapter gives us a belief set with some attractive properties: chiefly, for our purposes, interpersonal permissivism for full beliefs, even given objective Bayesianism. We have to set this up against the other properties, laid out in this chapter: intrapersonally permissivism cases and so on. I've argued that these latter properties do not threaten the approach.

Chapter 5

Conclusions

Let's recapitulate. We started with a discussion of the idea of permissivism, from the point of view of its most fervent and earliest critics. Richard Feldman, we saw (§1.1), focused his criticism on permissivism as it would play out in cases of peer disagreement, what he refers to as cases of supposed "reasonable disagreement." The primary way he thinks these cases might seem to arise is through agents somehow taking the evidence to point in different directions, or pointing with different strength. We analogized this to how Bayesians consider evidence to update one's priors, and how inferences to the best explanation might differ in a single case depending on which theoretical virtues the reasoner weights heavier. With this framing in mind Feldman's arguments against permissivism, namely that there should be a single rational way to consider the evidence, lest we end up with a deeply arbitrary rationality, seems less than convincing.

Roger White (§1.2) moves the discussion away from a focus on disagreement as activity, to disagreement as difference in doxastic attitude towards a proposition. His overarching worry, like Feldman's, was about arbitrariness, and his chief argumentative "move" reflects this. He claims that, were permissivism true, agents in permissive situations would be rationally permitted to change their beliefs without obtaining any new evidence; that is, they can "belief-toggle." In response, I proposed that White was insufficiently careful with the distinction between intra-personal and inter-personal permissivism; that is, whether permissivism holds for a single agent or between different agents. Consideration of the difference between these defangs the argument from White: the belief-toggling objection only works if every permissive case is intrapersonally permissive.

We then considered some positive arguments for permissivism from the literature. Miriam Schoenfield, we saw (§2.1), fully embraces the idea that different epistemic standards (understood as a subjective rule for moving from a body of evidence to a set of beliefs) can lead different agents to believe different things based on the same evidence. Indeed, it seems she embraces it *too* fully: there appears to be a tension in her theory between a willingness to step back from her current standards and consider other standards rational, and a denial that we can step back, consider other standards rational, and *therefore* be moved to adopt those standards. She seems, that is, unreasonably open to objections of the belief-toggling type.

Titelbaum and Kopec do better on this count (§2.2); they powerfully dispel the most pressing worries and present a positive case centred around a specific permissive case: the Reasoning Room. While obviously operating with some notion of epistemic standards analogous to Schoenfield's, their approach seems much more robust. The take-away from the first two chapters is that however worrying we might find the initial arguments from White and Feldman, the informal literature has plausible permissivist answers.

Chapters three and four presented, considered, and ultimately defended conditionally a positive proposal for a permissive view of full rational belief. The approach is based on the observation from William James that the epistemic value of believing truths and the epistemic disvalue believing falsehoods are in principle separable, and that different people might consider these values differently. We use this idea to present a model of the relation between full and partial credences, on the assumption that the latter represents our evidence. The “epistemic risk” model of rational doxastic states demands of epistemic agents that their full beliefs cohere in a certain way with their partial beliefs and their attitudes to the relative values of truth and falsehoods, that is, attitudes to epistemic risk. From these attitudes we develop an idea of epistemic utility, and the maximising of, in particular, *expected* epistemic utility. If different agents are allowed different attitudes to epistemic risk, the model will sometimes hold different attitudes to maximise expected epistemic utility, which in turn will demand different full beliefs of agents with identical credences.

There are a couple of conditional points for these conclusions worth noting more explicitly, as well as some possible objections. Firstly, it is not a given that there are different epistemically legitimate attitudes to epistemic risk. If there aren't (though I believe there are, §§3.3-4) the approach laid out will not deliver permissivism in the assumed circumstances. Secondly, it

is not a given that the permissivism we've gotten is actually the thesis we set out to defend. The latter was the following: sometimes, given a single body of evidence, different agents are allowed to rationally hold different doxastic attitudes to some proposition. What the approach gives us is the following: sometimes, given a single credence function, different agents with different attitudes to epistemic risk are allowed to hold different expected epistemic utility maximising attitudes toward some proposition. These are not obviously the same. I do believe (and have argued), however, that the idea of rational belief is captured by expected epistemic utility maximising belief, and that sometimes, the same body of evidence will mandate the same credences of different agents.

By way of possible objections, we noted that some of the points raised by White and Feldman apply in some capacity to the version of permissivism developed. In some cases, credence at certain thresholds, we get intrapersonal permissivism, and cannot in these cases rule out belief-toggling (§4.1). We fail to, in principle, rule out belief-toggling via change of standards (§4.2). And we remain somewhat unclear about how to describe an agent who reaches a different conclusion about some proposition than myself because she has a different attitude to epistemic risk (§4.3): rational? reasonable? understandable? None of these objections, for reasons explained, strike me as disqualifying, though it is important to recognize the problems they focus on as genuine consequences of the approach.

Whatever the take-away, I believe this has been a fruitful attempt at employing formal, precise work to answer a problem in more traditional epistemology. As has been made clear, not everything is given in using such methods: we have to enter into something like reflective equilibrium between the formal results and the more philosophical ideas. The last chapter was an attempt at such a balancing of considerations.

Bibliography

- Buchak, Lara. 2014. "Belief, Credence, and Norms." *Philosophical Studies* 169 (2): 1–27.
- Carnap, Rudolf. 1950. *Logical Foundations of Probability*. Chicago, IL, USA: Chicago University of Chicago Press.
- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2): 187–217.
- Clarke, Roger. 2013. "Belief Is Credence One (in Context)." *Philosophers' Imprint* 13:1–18.
- Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford University Press.
- Dorst, Kevin. 2019. "Lockeans Maximize Expected Accuracy." *Mind* 128 (509): 175–211.
- Douven, Igor. 2009. "UNIQUENESS REVISITED." *American Philosophical Quarterly* 46 (4): 347–361.
- Douven, Igor, and Hans Rott. 2018. "From Probabilities to Categorical Beliefs: Going Beyond Toy Models." *Journal of Logic and Computation* 28 (6): 1099–1124.
- Easwaran, Kenny. 2016. "Dr. Truthlove Or: How I Learned to Stop Worrying and Love Bayesian Probabilities." *Noûs* 50 (4): 816–853.
- . 2017. "The Tripartite Role of Belief: Evidence, Truth, and Action." *Res Philosophica* 94 (2): 189–206.

-
- Easwaran, Kenny, Luke Fenton-Glynn, Christopher Hitchcock, and Joel D. Velasco. 2016. "Updating on the Credences of Others: Disagreement, Agreement, and Synergy." *Philosophers'Imprint* 16:1–39.
- Feldman, Richard. 2007. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise Antony, 194–214. Oxford University Press.
- Foley, Richard. 1983. "Epistemic Conservatism." *Philosophical Studies* 43 (2): 165–182.
- Greco, Daniel, and Brian Hedden. 2016. "Uniqueness and Metaepistemology." *Journal of Philosophy* 113 (8): 365–395.
- Hjortland, Ole T., and Pål Antonsen. 2016. "Anti-Exceptionalism About Logic; research project at the University of Bergen," https://www.uib.no/sites/w3.uib.no/files/attachments/antiexceptionalism_project_short.pdf.
- Horowitz, Sophie. 2017. "Epistemic Value and the Jamesian Goals." In *Epistemic Consequentialism*, edited by Jeffrey Dunn Kristoffer Ahlstrom-Vij. Oxford University Press.
- Jackson, Elizabeth. 2021. "A Defense of Intrapersonal Belief Permissivism." *Episteme* 18 (2): 313–327.
- James, William. 1897. "The Will to Believe." In *The Will to Believe, and other essays in popular philosophy*. Longmans Green.
- Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4): 575–603.
- Kaplan, Mark. 1996. *Decision Theory as Philosophy*. Cambridge University Press.
- Kelly, Thomas. 2014. "Evidence Can Be Permissive." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri, 298–312. Wiley Blackwell.
- Leitgeb, Hannes. 2017. *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.
- Lyon, Aidan. 2010. "Philosophy of Probabililty." In *Philosophies of the Sciences: A Guide*, edited by Fritz Allhoff. Wiley-Blackwell.

-
- Meacham, Christopher. 2013. "Impermissive Bayesianism." *Erkenntnis* 79 (Suppl 6): 1185–1217.
- Meyer, Albert R., and Ronitt Rubinfeld. 2005. "Missed expectations? Course notes for Mathematics for Computer Science, MIT," <https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-042j-mathematics-for-computer-science-fall-2005/readings/ln14.pdf>.
- Pettigrew, Richard. 2013. "Epistemic Utility and Norms for Credences." *Philosophy Compass* 8 (10): 897–908.
- . 2016. *Accuracy and the Laws of Credence*. Oxford University Press UK.
- . 2017. "Epistemic Utility and the Normativity of Logic." *Logos and Episteme* 8 (4): 455–492.
- . 2021. "On the Pragmatic and Epistemic Virtues of Inference to the Best Explanation." *Synthese* 199 (5-6): 12407–12438.
- . ms. "Epistemic Risk and the Demands of Rationality." Downloaded from: <<https://philpapers.org/archive/PETERA-4.pdf>>.
- Rosen, Gideon. 2001. "Nominalism, Naturalism, Epistemic Relativism." *Noûs* 35 (s15): 69–91.
- Schoenfield, Miriam. 2014. "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Noûs* 48 (2): 193–218.
- . 2019. "Permissivism and the Value of Rationality: A Challenge to the Uniqueness Thesis." *Philosophy and Phenomenological Research* 99 (2): 286–297.
- Smith, Julia Jael. 2020. "Unacknowledged Permissivism." *Pacific Philosophical Quarterly* 101 (1): 158–183.
- Steinberger, Florian. 2019. "Accuracy and Epistemic Conservatism." *Analysis* 79 (4): 658–669.
- Titelbaum, Michael G. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press.

Titelbaum, Michael G., and Matthew Kopec. 2019. “When Rational Reasoners Reason Differently.” In *Reasoning: New Essays on Theoretical and Practical Thinking*, edited by Magdalena Balcerak Jackson and Brendan Balcerak Jackson, 205–231. Oxford University Press.

———. ms. “Plausible Permissivism.” Downloaded from: <<https://philpapers.org/archive/TITPP-2.pdf>>.

Wedgwood, Ralph. 2002. “The Aim of Belief.” *Philosophical Perspectives* 16:267–97.

———. 2017. “Epistemic Teleology: Synchronic and Diachronic.” In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeff Dunn, 85–112. Oxford: Oxford University Press.

Weisberg, Jonathan. 2020a. “Belief in Psyontology.” *Philosophers’ Imprint* 20 (11).

———. 2020b. “Could’ve Thought Otherwise.” *Philosophers’ Imprint* 20 (12).

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

———. 2010. “Evidential Symmetry and Mushy Credence.” In *Oxford Studies in Epistemology*, edited by T. Szabo Gendler and J. Hawthorne, 161–186. Oxford University Press.

———. 2014. “Evidence Cannot Be Permissive.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri, 312–323. Wiley Blackwell.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.