

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/126566>

Copyright and reuse:

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Political Liberalism and Epistemic Permissivism

by

Muralidharan Anantharaman

A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy in Philosophy

University of Warwick, Department of Philosophy

June 2019

Contents

Acknowledgments	iii
Declaration	iv
Abstract	v
Chapter 1: Introduction	1
1.1 The Standard Account of Political Liberalism	
1.1.1 The Reasonable and the Rational	
1.1.2 The “Justifiable To” Relation and Epistemic Access	
1.1.3 Reasonable Disagreement and the Consensus and Convergence Models of Public justification	
1.2 Rational vs Blameless Disagreement between Reasonable Persons	
1.3 Permissivism	
1.4 Chapter Summaries	
Chapter 2: The Burdens of Judgement and Permissivism	27
2.1 Evidence	
2.2 The First Burden	
2.3 The Second Burden	
2.4 The Third Burden	
2.5 The Fourth Burden	
2.6 The Fifth Burden	
2.7 The Sixth Burden	
Chapter 3: The Best Version of Permissivism	46
3.1 The Evidence Pointing Problem	
3.2 The Arbitrariness Objection	
3.3 Epistemic Standards	
3.4 On the Irrelevance of Procedural Considerations to Permissivism	
3.5 Your Belief Makes it Real	
Chapter 4: Intrapersonal Uniqueness about Epistemic Standards	65
4.1 Extending the Arbitrariness Objection	
4.1.1 Choosing for No Reason	
4.1.2 Choosing for Practical Reasons	
4.1.3 Choosing for Epistemic Reasons	

4.1.4 Viewing Alternative Standards	
4.2 Extending the Evidence Pointing Problem	
Chapter 5: Permissivism, Cognitive Capacities and Limitations	92
5.1 Truth Conduciveness Relative to Cognitive Capacities	
5.2 Difficult Questions and Cognitive Limitations	
5.3 Limitation of Epistemic Obligations by Cognitive Capacity	
5.4 The Naïve Account	
5.5 Forgiving Accounts	
Chapter 6: Indexing Epistemic Standards to Agent's Epistemic Values	110
6.1 Instrumentalist Permissivism	
6.2 Deciding Without Presupposing a Likelihood	
6.3 Pragmatic Goals	
Chapter 7: Epistemic Conservatism and Permissivism	130
7.1 Epistemic Conservatism	
7.2 Motivating Conservatism	
7.3 Objections to Conservatism	
7.4 Weakening Conservatism	
7.4.1 Differential Conservatism	
7.4.2 Defeating Conditions	
7.4.3 Mere Beliefs	
7.4.4 Dynamic Conservatism	
Chapter 8: Conclusion	156
Bibliography	159

Acknowledgments

I would like to thank my super Professor Fabienne Peter for her guidance and support in the writing of this thesis. I would also like to thank my family: my parents, my siblings and my wife Nandini for their moral and emotional support during this period. I could not have achieved this without any of them.

Declaration

I, Muralidharan Anantharaman, hereby, declare that the following thesis is my own work. I also confirm that this thesis has not been submitted for a degree at any other university.

Muralidharan Anantharaman

17/06/2019

Abstract

The supposed fact of reasonable disagreement plays a crucial role in standard accounts of political liberalism. The standard account posits reasonable disagreement as arising primarily from the supposed fact that, in at least some circumstances, people can permissibly respond in different ways to the same evidence. That is to say, the standard account presupposes permissivism: It is possible that more than one doxastic attitude towards a proposition is rationally permissible, given a body of evidence. In this thesis, I shall show that the standard account is mistaken in presupposing permissivism by arguing for Uniqueness: Given a total body of evidence at most one doxastic attitude is rationally permissible. I shall argue for this by first rehearsing White's two objections to permissivism and showing that, at the very least, they show that an intrapersonal version of Uniqueness is true. That is to say, at most one doxastic attitude is justified for a given agent given a total body of evidence. I shall then present different accounts of permissivism consistent with intrapersonal uniqueness and show that none of these accounts are successful. If my argument is successful, the standard account of political liberalism may have to be revised.

Chapter 1: Introduction

The standard version of political liberalism consists of two claims. The first claim is that there is a public justification requirement according to which the exercise of political authority is legitimate only if it is publicly justifiable. The second claim is that only liberal political authority is publicly justifiable since it is the only type of political authority that does not appeal to views, such as sectarian religious or moral doctrines, that would be subject to reasonable disagreement. Epistemic permissivism is the claim that people with the same total body of evidence can rationally disagree. In this thesis, I shall argue firstly that the standard version of political liberalism presupposes epistemic permissivism and secondly that permissivism is false. A consequence of this is that certain important moral claims that political liberals standardly take to be subject to reasonable disagreement are not. If so, then the standard account of political liberalism is in need of revision. In order to evaluate these two claims, I shall, in this chapter, explain what I mean by the terms “political liberalism” and “epistemic permissivism”.

1.1 The Standard Account of Political Liberalism

We can more easily understand the term political liberalism if we examine what I take to be a core claim: Laws based on controversial moral or religious doctrines are illegitimate because such doctrines are subject to reasonable disagreement. Implicit in this claim is a public justification requirement according to which:

Public Justification Requirement: The exercise of political authority is legitimate only if it is publicly justifiable.

I shall first clarify a few terms. By the exercise of political authority, I refer to the action of a state. Such state actions include coercive laws and policies, like those which prohibit the construction of mosques in a city, those which prohibit murder or those which set a minimum wage, and also to non-coercive laws like marriage and contract laws¹. Similarly, pronouncements made by political leaders may also be covered by this principle. For instance, if a public official were to declare that certain religious beliefs were necessary for good citizenship, political liberals would

¹ HLA Hart (1964) has noted that contract and marriage law are non-coercive. Violation of these laws only results in the contract or marriage being nullified. However, according to Hart, nullity is not a sanction. While sanctions are typically undesirable and it is the threat of this undesirable consequence which creates the coercive character of some laws, nullity can often be seen as neutral or even desirable. Even if laws which define marriage as only between a man and a woman are not themselves coercive, political liberals may nevertheless think they are illegitimate if they are not publicly justifiable.

consider that statement illegitimate since we generally regard it as not being publicly justifiable in a pluralistic society.

There are other versions of the public justification requirement in which the justificandum is different. For instance, according to Rawls's Liberal Principle of Legitimacy (Rawls 1993: 137), only the constitution of a polity needs to be publicly justifiable. The exercise of political power is legitimate only if it is in accordance with such a constitution. On this version of the principle, the constitution may prescribe a procedure which is publicly justifiable, but that procedure need not pick out laws or policies which themselves are publicly justifiable. To take another example, Gaus's Deliberative Public Justification Principle requires what he calls "bona fide rule[s] of social morality" (Gaus 2011: 27) to be publicly justifiable. While the public justification requirement, as I have specified it, takes laws or policies as the justificandum, the question of exactly what the justificandum for public justification is, does not matter. What is important is that public justification of the justificandum depends on the epistemic justification of certain relevant beliefs. While, I shall explain how this can be in the case where laws and policies are the relevant justificanda, this should apply equally well to any other plausible candidate justificandum.

By publicly justifiable is meant the following:

Basic Public Justification: The exercise of political authority is publicly justifiable only if it is justifiable to each reasonable person subject to it.

There are two key terms which require clarification, "justifiable to" and "reasonable person". With regards to the core political liberal claim, little has been said, in the literature, about what reasonable disagreement amounts to. I shall, later in this chapter, defend an interpretation of "reasonable disagreement" as "rational disagreement between reasonable persons". This, as I shall argue, is the most defensible interpretation of the requirement that each reasonable person have epistemic access to the reasons that justify a law². In what follows, I shall first clarify the term "reasonable person" before addressing how "justifiable to" relates to "reasonable disagreement".

² To be clear, the claim that reasonable disagreement is rational disagreement between reasonable persons is not part of the standard account of Political Liberalism.

1.1.1 The Reasonable and the Rational

In the standard version, “Reasonable person” is a term that demarcates the constituency to whom justification is owed. The fact that a given law is not justifiable to unreasonable persons does not undermine the legitimacy of the law. This justificatory constituency has been defined in various ways. For instance, according to Rawls and Jonathan Quong, a reasonable person is willing to, firstly, propose and honour fair terms of cooperation. Secondly, a reasonable person is able to recognise what Rawls calls the “burdens of judgment” and to recognise their consequences (Rawls 1993: 49; Quong 2010: 37-38). I shall explain what the burdens of judgment are later in the discussion on rational disagreement. There are, of course, alternative accounts of a reasonable person. These accounts add or subtract from the content of reasonableness, thus, respectively narrowing or widening the constituency of justification. While there are important issues raised by the choice of the conception of reasonableness³, the question of exactly which conception of reasonableness is appropriate is not of concern in this thesis. All that is needed is to take note of the claim that reasonableness picks out a constituency to whom justification is owed. It is justification among this constituency which matters for legitimacy. Justification of laws to persons who are not being reasonable is irrelevant with regards to satisfying the public justification requirement.

To clarify this point, it may seem as if there is an assumption here that some people are always reasonable and some people are always unreasonable. However, this is not the case: justification is addressed to persons when they are being reasonable. If, for instance, being reasonable is about being willing to propose and abide by fair terms of cooperation provided that others are as well, then an agent is part of that constituency, insofar as the demands she makes on others are consistent with this willingness. Suppose she is currently not willing to propose and abide by fair terms of cooperation because she believes that laws should conform to her account of religious truth: Her demand that she be permitted to construct a house of worship in a given plot of land could still be reasonable because she could consistently make such a demand even if she were willing to propose and abide by fair terms of

³ See for instance the argument advanced by Estlund (1998) that an over-inclusive conception of reasonableness opens up the public justification requirement to a self-defeat objection because the public justification requirement would itself not be justifiable to all reasonable persons. On the other side of the issue is Bajaj (2017) who argues that some under-inclusive conceptions of reasonableness, where being reasonable involves accepting the public justification requirement, risk being ad-hoc.

cooperation. In making this demand, we can treat her as being reasonable. A law forbidding the building of houses of worship of her religious denomination would be illegitimate, because it would not be justifiable to her even if she were willing to propose and abide by fair terms of cooperation. By contrast, suppose she also were to demand that people of other religions not practice theirs. She could not consistently make this demand if she were to be willing to propose and abide by fair terms of cooperation. Therefore, the fact that laws guaranteeing religious freedom are not justifiable to her does not make them illegitimate. This illustrates the way in which the account of reasonable persons is an idealisation. The constituency of justification is not people as they actually are, but people as they would be if they were being reasonable (Larmore 2008; Quong 2010).

In the standard version, the word “reasonable” has been differentiated from the word “rational”. For Rawls 1993: 48-51), being rational pertains to an agent exercising her capacities to a) weigh her final ends and projects against one another and revise them in light of their importance to her as well as b) to choose and pursue the various means that will achieve said ends. Rawls distinguishes being rational from being reasonable, which as I have mentioned, refers primarily to a disposition to engage in fair social cooperation (ibid). In this thesis, my use of the term “rational” is purely epistemic and hence is orthogonal to Rawls’s distinction between the rational and the reasonable. In particular, by a rational belief, I mean one which is supported by the agent’s evidence. Correspondingly, a rational disagreement is one in which both disagreeing parties’ respective beliefs are rational. For instance, a religious extremist might be rational in Rawls’s sense of the term by being able to pursue the means to her ends and acting on the basis of the values she accepts. However, insofar as she is willing to pursue violent means to her ends, she is not reasonable. Moreover, insofar as her extreme beliefs do not fit her evidence, she is not rational in the way I shall be using the term in this thesis. This is not to say that “reasonable”, especially as applied to disagreement, is a purely epistemic notion, rather it is a political idea which contains epistemological elements (ibid: 62). While political liberals have in general been vague about what these epistemological elements are, I will later defend the claim that the epistemic component of reasonableness refers to epistemic rationality. That is to say, reasonable disagreement is rational disagreement between reasonable persons.

1.1.2 The “Justifiable To” Relation and Epistemic Access

To clarify what is meant by “justifiable to”, a law *L* being justifiable to an agent Alice involves two senses of justification. The first is the law being justified by practical reasons. For instance, one may think that what justifies laws against murder is that murder is wrong or that what justifies a social safety net is that such a policy would stave off the worst effects of poverty. Associated with these practical reasons are propositions, which can be moral or non-moral. For instance, the claims that murder is wrong or that we ought to alleviate the worst effects of poverty are moral propositions. The claim that a guaranteed minimum income changes the distribution of wealth in certain ways is a non-moral proposition. Both types of propositions can be associated with reasons to enact, abolish or enforce certain laws and policies.

The second sense of justification is epistemic: In order for *L* to be justifiable to Alice, the practical reasons which justify *L* must be epistemically accessible to her. Little has been said about what this access condition involves; Larmore’s (2015) and Gaus’s (2011) accounts of political liberalism suggest that, at the very least, Alice having access to a reason *R* involves her being epistemically justified in believing that *P*, the proposition associated with *R*. More recently, Peter (2018) has argued that the propositions associated with the reasons which justify political decisions should be justifiably believed by each reasonable person. Later in this chapter, I shall defend the claim that being epistemically justified in believing that *P* is necessary and sufficient in having access to the reason, *R*. For now, it is sufficient to note that an agent must have some level of epistemic access to the practical reasons that justify a law, in order for that law to be justifiable to her.

1.1.3 Reasonable Disagreement and the Consensus and Convergence Models of Public justification

Broadly speaking, there are two ways in which laws can come to be publicly justifiable: a consensus model and a convergence model. On the consensus model⁴, there is some set of practical reasons *R* which both justifies a law *L* and which is accessible to all reasonable persons. By contrast, proponents of the convergence model (Gaus 2011; Vallier 2011; Bajaj 2017) argue that on any sufficiently inclusive account of reasonableness, the set of reasons that are accessible to all reasonable persons is too small to justify any law. On the convergence model, they propose that laws tend to be publicly justifiable because

⁴ See Rawls (1993), Estlund (1998), Quong (2010) and Larmore (2008; 2015) and as proponents of this view.

they can be justified by more than one set of reasons $R_1, R_2 \dots R_n$ and at least one or other of R_1 to R_n is accessible to any given reasonable person. For instance, a law preventing one person from enslaving another could be justified because slavery produces lots of misery, because slavery involves the subordination of one person to another, or because it violates a natural right to liberty. On the convergence view, at least one or other, but not necessarily the same one, of these reasons is epistemically accessible to all reasonable persons.

Despite the disagreements between proponents of consensus and convergence models of public justification, they will both agree on how laws fail to be publicly justifiable. The central instances of a failure of public justifiability involve cases in which there are certain reasons R_1 to R_k which justify a law L , but R_1 to R_k are not accessible to all reasonable persons. Instead, some other reasons R_{k+1} to R_n are accessible to some reasonable persons. However, R_{k+1} to R_n do not, either individually or in conjunction, justify L . Hence, on any plausible model of public justification, laws fail to be publicly justifiable because not all reasonable persons have epistemic access to the reasons that can justify those laws.

It thus makes sense to think of reasonable disagreement about a proposition as disagreement between reasonable persons that arises from a lack of access to that proposition. For instance, if, for a reason R and the associated proposition P , having epistemic access to R means being epistemically justified in believing that P , reasonable disagreement about P refers to epistemically justified or rational⁵ disagreement about P . Provided that lacking epistemic justification for P is sufficient to deprive an agent of access to R , if two reasonable persons rationally disagree about P , at least one of them lacks access to R .

1.1.4 Reasonable Disagreement and the Burdens of Judgment

Summing up so far, I have described how, on the public justification requirement, publicly justifiable laws can be justified on the basis of practical reasons which are epistemically accessible to all reasonable persons. Failures of public justification occur because one or more reasonable persons are unable to access these reasons. Given the role reasonable disagreement is supposed to play in undermining public justification, I have stipulated that reasonable disagreement is disagreement that arises from a lack of epistemic access to these reasons. Approaching this from the other direction, we might ask how reasonable

⁵ I shall use epistemically justified and rational interchangeably.

disagreement might arise. According to Rawls, people reasonably disagree because of what he calls the burdens of judgment. The burdens of judgment describe the “many hazards involved in the correct (and conscientious) exercise of our powers of reason and judgment in the ordinary course of political life” (Rawls 1993: 55-56). The aim of any satisfactory account of these burdens is to show that disagreement caused by these burdens does not impugn the reasonableness of the disagreeing parties. Rawls provides an admittedly non-exhaustive list of these burdens: (ibid: 56)

“a. The evidence – empirical and scientific – bearing on the case is conflicting and complex and thus hard to assess and evaluate.

b. Even where we agree fully about the kinds of considerations that are relevant, we may disagree about their weight, and so arrive at different judgments.

c. To some extent, all our concepts, and not only moral and political concepts, are vague and subject to hard cases; and this indeterminacy means that we must rely on judgment and interpretation (and on judgments about interpretations) within some range (not sharply specifiable) where reasonable persons may differ.

d. To some extent (how great we cannot tell), the way we assess evidence and weigh moral and political values is shaped by our total experience, our whole life up to now; and our total experiences must always differ. Thus, in a modern society with its numerous offices and positions, its various divisions of labor, its many social groups and their ethnic variety, citizens’ total experiences are disparate enough for their judgments to diverge, at least to some degree, on many, if not most, cases of any significant complexity.

e. Often, there are different kinds of normative considerations of different force on both sides of an issue and it is difficult to make an overall assessment.

f. Finally, as we note in referring to Berlin’s view (V:6.2), any system of social institutions is limited in the values it can admit so that some selection must be made from the full range of moral and political values that might be realised. This is because any system of

institutions has, as it were, a limited social space. In being forced to select among cherished values, or when we hold to several and must restrict each in view of the requirements of the others, we face great difficulties in setting priorities and making adjustments. Many hard decisions may even seem to have no clear answer.” (ibid: 56-57)

These listed burdens are supposed to represent the more obvious sources of reasonable disagreement. Other political liberals either endorse Rawls’s list or offer a similar account. For instance, Larmore’s list includes, with reference to the first and fifth burdens, the complexity of an issue; as with the fourth burden, different background experiences; and, as with the sixth burden, the fact that we are sometimes forced to choose one value over another as reasons why reasonable people can rationally disagree (Larmore 2015: 69). Gaus’s account of rational disagreement invokes the fact that we have limited cognitive resources as a source of reasonable disagreement (2011). As I shall argue in Chapter 5 this relates to the complexity of many questions in the following way: More complex questions require more cognitive resources from an agent in order for her to avoid making mistakes. Since people have limited cognitive resources, questions whose complexity exceeds the available cognitive resources will be subject to blameless errors in reasoning. Reasonable disagreement arises from such blameless errors. Thus Gaus’s account also posits the complexity of an issue as being a cause of reasonable disagreement.

Summing up, I have presented what I take to be the standard account of political liberalism. On the standard account, illiberal laws are illegitimate because they cannot be publicly justified as they are based on sectarian moral or religious doctrines: propositions which are subject to reasonable disagreement. Reasonable disagreement about these doctrines makes it the case that not all reasonable citizens have epistemic access to the reasons that would justify these illiberal laws. I have also presented how political liberals tend to account for reasonable disagreement: the burdens of judgment.

1.2 Rational vs Blameless Disagreement between Reasonable Persons

In my exposition of the standard account of political liberalism, I have left exactly what is meant by “epistemic access” and “reasonable” unspecified. This is primarily because political liberals have been vague in their accounts of what is meant by reasonableness.

For instance, this ambiguity can be found in the way Rawls and many others following him have presented the idea of reasonable disagreement: The exercise of human reason under conditions of freedom is supposed to lead inevitably to widespread disagreement about moral, religious and political doctrines⁶ (Rawls 1993: xxvi; Larmore 2008: 140; Vallier 2011: 261; Gaus 2011: 2). There are many ways in which we might interpret the phrase “exercise of human reason”. On a narrow interpretation, an agent exercises her capacity to reason only if she reasons in a way that would lead her to being epistemically justified in the attitude she would end up with once she finished reasoning. On this narrow interpretation, the disagreement produced by the exercise of human reason is rational disagreement. There is a wider interpretation according to which an agent exercises her capacity to reason so long as she is sincere and conscientious in her effort to find the truth. As long as both disagreeing agents are sincere and conscientious⁷, then their disagreement is blameless⁸.

Whatever “reasonable” means, it must have some epistemic component. My concern in this section is to defend the view according to which “reasonable” entails epistemically rational. As mentioned before, in order for a law or policy to be justifiable to a person, the practical reasons R which justify that law must be epistemically accessible to her. Reasonable disagreement is disagreement which arises from some reasonable persons lacking epistemic access to R. Therefore, the

⁶ Plausibly, accepting a doctrine need not just be a matter of believing certain propositions. For instance, it may also involve adopting certain sorts of practical stances or outlooks. Nevertheless, I shall take it that on any plausible account of what a doctrine is, believing or being prepared to believe some set of core propositions is a necessary part of accepting a doctrine. My concern in this thesis is purely with this propositional aspect of doctrines.

⁷ Certainly, there are other sources of this ambiguity. As Freeman notes (2007: 345-346), Rawls claims that the reasonableness is not an epistemological idea, but a political one with epistemological elements (1993: 62). The question that I am attempting to address in this section is what exactly this epistemological component consists in. Nothing that Rawls says rules out the possibility that the epistemic component of reasonableness just is epistemic rationality.

⁸ See, for instance Quong (2010: 295-296) who supposes that reasonable doctrines need not be well supported by the evidence. Quong’s interpretation of reasonableness seems to involve a wider interpretation of what it means to exercise the capacity for reasoning. See also Jonch-Clausen and Kappel (2015: 377-378) who take beliefs to be rational, in a certain sense, if the agent believes that those beliefs are supported by her evidence. While Jonch-Clausen and Kappel use the term rational to describe said beliefs this is not the same sense of rationality I use in this paper. Rationality, as I use it refers to what the evidence actually supports and not merely to what the agent believes the evidence supports. After all, it seems all too possible to have false or badly formed beliefs about what the evidence supports. Moreover, true beliefs that are formed on the basis of what the evidence is merely believed to support, even when combined with some suitable external condition like safety could not be knowledge. Therefore Jonch-Clausen and Kappel’s account of rationality is more akin to blameless belief than rational belief.

question of how to define reasonable disagreement depends on our conception of epistemic access and vice versa. Let me explain.

Suppose that the following account of access is correct.

Rational: A practical reason R is epistemically accessible to an agent if and only if it would be epistemically rational⁹ for her to believe the associated proposition P.

According to Rational, for someone to lack access to R, it would have to be the case that it is not rational for her to believe that P. If it were rational for her to believe that P, she would have access to R. If, instead¹⁰, it was rational for the agent to suspend judgment about P or disbelieve P, then she would not have access to R. Since reasonable disagreement is disagreement that arises from some people lacking epistemic access to R, if Rational is the correct account of access, then it is reasonable to suspend judgment about P or disbelieve P only if it is rational to have, respectively, either of those doxastic attitudes towards P. Since a disagreement is reasonable only if it would be reasonable for each of the disagreeing parties to have their respective beliefs, disagreement is reasonable only if it is rational.

There is an inverse relationship between the account of access and the account of reasonableness. If epistemic access needs to be very robust, it is easier to fall short of such access and hence easier for disagreement to count as reasonable. To see why this is the case, consider an account of access that is more robust¹¹ than Rational.

Blameless: A practical reason R is epistemically accessible to an agent if and only if a) it would be epistemically rational for her to

⁹ In this thesis, I shall be using the terms epistemically rational and epistemically justified interchangeably. More specifically, I do not distinguish between rationality, warrant, entitlement and justification. By rational or justified, I refer to whatever property a true belief must have in addition to some external condition like safety in order for it to count as knowledge.

¹⁰ For simplicity of presentation, I am assuming that if one doxastic attitude is rational for an agent, no other attitude is rational for her. The argument presented in this section does not depend on this assumption. Furthermore, I shall defend this assumption later in Chapter 3.

¹¹ While weaker accounts of epistemic access are logically possible, they do not seem particularly plausible and therefore are not salient. Suppose an agent had epistemic access to R whenever she had at least some evidence for P. Then even if she ought to believe not-P because her total body of evidence favoured not-P, a law that was justified by R would be justifiable to her. This generates the wrong result, whereby, if for a given question of whether P, there is evidence on both sides of the issue, then both R and not-R are accessible to the agent. If anything, when there is evidence on both sides of the issue, it should render P subject to reasonable disagreement. Accounts of access weaker than Rational will inevitably reach the wrong result in at least some cases.

believe that the associated proposition P and b) she could only fail to believe that P by reasoning so badly as to be blameworthy in her reasoning.

Blameless is a more robust account of epistemic access than Rational. If Rational is true, then if an agent is rational in believing that P she necessarily has access to R. However, this need not be the case if Blameless was true instead. Consider the following case

Immigration: One reason, R, for relaxing immigration restrictions is that it would help migrants without imposing significant costs on natives and we should help others if we can do so with only negligible cost to ourselves. Associated with this practical reason, R are two propositions. One is a normative proposition which claims that we ought to help others if we can do so at little cost. The second proposition is that relaxing immigration restrictions would help migrants without imposing significant costs on natives. Suppose it is rational for an agent, Alice to believe that relaxing immigration restrictions would not cause native workers to lose their jobs. It would be rational for Alice to believe this because it is supported by her evidence. However, the evidence is very complicated and her cognitive abilities are limited. Therefore, she could also blamelessly but not rationally believe that relaxing immigration restrictions would cause native workers to lose their jobs.

In this situation, according to Blameless, Alice lacks epistemic access to R. Since reasonable disagreement is disagreement that arises from lacking access to that reason, if Alice blamelessly believed that relaxing immigration restrictions would cause natives to lose their jobs, her belief would be reasonable. By contrast, if Rational was the correct account of access, then since Alice has epistemic access to R, her blameless belief that relaxing restrictions would cause natives to lose their jobs would not be reasonable.

On the other hand, if it was rational for Alice to believe that relaxing immigration restrictions would cause natives to lose their jobs, then under both Rational and Blameless, she lacks access to R. Therefore, on both accounts of access her belief would be reasonable. This should make clear that a more robust account of epistemic access corresponds to a weaker account of reasonable disagreement. Disagreements that count as reasonable under Rational would also count as

reasonable under Blameless. However the reverse is not necessarily true. Correspondingly, we should expect that cases in which an agent has access to the reason under Blameless should also count as ones in which she has access under Rational. Suppose an agent, Betty's, evidence supports the proposition that relaxing immigration restrictions would not cause locals to lose jobs. Suppose further that her evidential situation is such that she could not blamelessly believe otherwise. Then, even according to Blameless, she has epistemic access to R. It follows trivially that since it is rational for her to believe that relaxing restrictions would not cause locals to lose jobs, she also has epistemic access under Rational.

Summing up, there is an inverse relationship between the account of epistemic access and the account of reasonableness. As I shall explain later on in this chapter, permissivism is the claim that rational disagreement is possible even when people share the same total body of evidence. Permissivism is relevant to the question of reasonable disagreement only if reasonable disagreement entails epistemically rational disagreement. Therefore, as I have mentioned before, in this thesis, I shall defend a corresponding account of epistemic access in which being justified in believing that P is both necessary and sufficient for having epistemic access to the associated reason R. That is to say, the view I shall be defending in this section is Rational whereas the view that I shall be arguing against will be Blameless. As mentioned earlier, the account of reasonable disagreement that corresponds to Rational is rational disagreement between reasonable persons. Likewise, the account of reasonable disagreement that corresponds to Blameless is blameless disagreement between reasonable persons.

In setting up a distinction between blameless beliefs and rational beliefs, I need to explain how it is that an agent's beliefs fail to be rational, but can still be blameless. There are a number of ways this might come about. For instance, there might be some cases which, perhaps unknown to her, are more complicated and her usual amount of effort does not result in beliefs which fit her evidence in those cases. Alternatively, her cognitive capacities might be impaired by illness or old age in ways that make it inappropriate to blame her when she responds inappropriately to her evidence. Depending on the account of blameless belief, there could be significant variation in which beliefs and hence which disagreements count as blameless. Given that my argument should succeed regardless of the account of blameless belief, I shall not commit myself to any single account of blamelessness. Whatever the account of blameless belief, the account of reasonableness as blamelessness is logically weaker than the account of reasonableness as rationality.

The reason for this is that when someone's beliefs are rational, they are necessarily blameless as there is nothing to be blamed for. We can thus divide blameless beliefs into two types: Those which are both blameless and rational and those which are blameless but not rational. As such, the discussion will focus on merely blameless beliefs, beliefs which are blamelessly but not rationally held. Correspondingly, the discussion will focus on merely blameless disagreements, disagreements in which, both parties' beliefs are blamelessly held and at least one of the parties' beliefs is merely blamelessly held.

The question that concerns us is: which conception of reasonable disagreement is the most defensible? I shall argue that political liberals have decisive reason to take reasonable disagreement to be rational disagreement between reasonable persons. The three reasons for this claim are as follows: The first reason is that on any other account of reasonable disagreement, the normative requirement that people be epistemically reasonable becomes incoherent. The second reason is that on such weaker accounts of reasonable disagreement, far too many doctrines are subject to reasonable disagreement and this would make it impossible to publicly justify liberal laws. The third reason is that the normative significance of these weaker accounts of epistemic reasonableness is questionable. I shall hereby present each of these objections in turn.

The first objection to reasonable disagreement as epistemically blameless disagreement between reasonable persons is that on any more robust account of epistemic access like Blameless, the corresponding normative requirement that people be reasonable becomes incoherent. First, I shall explain what I mean by the normative requirement to be reasonable. When we are attempting to justify a law to others, and they fail to accept the reasons that we provide even though we think they ought to, we are inclined to demand of them that they be reasonable. This presupposes that there is some normative requirement to be reasonable such that when everyone is being reasonable, they will endorse laws and policies which will be acceptable to and accepted by others insofar as they too are being reasonable. The intuition here is that reasonable disagreement would not be normatively significant unless it was the case that people ought to be reasonable. Supposing that the public justification requirement therefore seems to have a correlative duty to be reasonable, this raises the further question as to what the content of this duty is.

As I have mentioned before, it is plausible that there is an epistemic aspect to being reasonable. For instance, if a climate change denialist were to object to some

sensible climate policy on the supposed grounds that, according to her, there is no anthropogenic global warming we might rightly regard her as unreasonable on the basis that denying climate change is sufficiently epistemically bad that her rejection of climate change policy cannot be counted as reasonable. We might similarly regard creationists, Holocaust deniers and other sorts of conspiracy theorists (Cassam 2016) as unreasonable too. The thought here is that apart from the other commitments associated with being reasonable, people also have an obligation to be epistemically reasonable in accepting moral and non-moral beliefs. It is these beliefs, subject to the other constraints of reasonableness, which people would accept if they were being epistemically reasonable which determine whether they have access to sufficient practical reasons to accept a law. If Rational is the correct account of access, being epistemically reasonable requires being epistemically rational. After all, it is only when all otherwise reasonable persons are epistemically rational, will they all accept all and only those laws which are justifiable to them. By contrast, if Blameless is the correct accounts of access, the normative requirements of reasonableness are weaker. Reasonable persons only need to have beliefs which are blamelessly formed in order to accept all and only those laws which are justifiable to them. My claim is that accounts of access like Blameless are objectionable because there is something incoherent about this weaker conception of reasonableness.

My claim relies on the intuition that normative requirements and permissions supervene on reasons. Satisfying these normative requirements involves responding adequately to these reasons (Parfit 2011; Scanlon 2013). Just as we can divide up our normative requirements into fact-relative, evidence-relative or belief-relative requirements, we can also divide up our reasons into fact-relative, evidence-relative or belief-relative reasons. The fact that the glass contains petrol and not whiskey is a fact-relative reason for Alice not to drink it. Correspondingly, Alice has a fact-relative normative requirement not to drink the liquid in the cup. Suppose, however, that she justifiably believes, on the basis of her evidence, that it is whiskey. Then, she has an evidence relative normative permission to drink the liquid in the cup. Each type of normative requirement or permission supervenes on the corresponding type of reasons. What is involved in being reasonable is responding to one's evidence-relative reasons. The weaker accounts of epistemic reasonableness are incoherent because they seem to issue conflicting demands. I shall explain why with the following.

Suppose Blameless is the correct account of epistemic access. Then, as long as someone's beliefs are blamelessly formed, they are epistemically reasonable. Consider the case in which there is some proposition P that is supposedly subject to reasonable disagreement only because some people blamelessly believe not-P even if it would have been rational for all of them to believe that P. If being blameless in your beliefs is sufficient to be epistemically reasonable, then, when everyone is being reasonable, some people would have a reasonable, but not rational belief that not-P. Since, by assumption, they are being reasonable, it would be inappropriate to demand of those that believe not-P that they should be reasonable.

However, this seems to conflict with the intuition that when people are not responding rationally to their evidence, it is appropriate to demand that they respond better. Suppose Alice has a merely blameless belief that not-P and that it would be rational for her to believe P given her evidence. In demanding that Alice respond properly to her evidence, we are not demanding that she comply with some fact-relative obligation that binds her regardless of her epistemic situation. The demand is that she should respond to the epistemic reasons that she has available to her. Therefore, the demand is evidence-relative and is binding on her *because* of her epistemic situation. Therefore, if merely blamelessly believing is sufficient for being epistemically reasonable, we seem to be both claiming that when she has a merely blameless belief, she has satisfied all the evidence-relative normative requirements that apply to her and that there are evidence-relative normative requirements she has failed to satisfy. This is inconsistent. Any similarly weak account of reasonableness will be subject to the same objection. On these accounts, one can simultaneously claim that the requirements of reasonableness have been satisfied and that there are further normative requirements that have not been satisfied. Given that these weaker accounts of reasonableness yield inconsistent claims about whether the normative requirements of reasonableness are satisfied, then the only account of reasonableness which does not is one whereby reasonable disagreement is rational disagreement between reasonable persons.

One way to respond to the above argument is to claim that there need not be any conflict. While it is appropriate to demand that people be rational when they are not, this demand is distinct from a demand that they be reasonable. Therefore, or so the response could go, while it would be inappropriate to demand that people be reasonable when they already are, demanding that they be epistemically rational is not inappropriate because being rational is a distinct requirement from being

reasonable. When someone is reasonable, it does not follow that she has satisfied all her evidence-relative normative requirements, only, that she has satisfied those that pertain to being reasonable. However, it is not incoherent to say that Alice has satisfied all the evidence-relative normative requirements that pertain to being reasonable, but has not satisfied other evidence-relative normative requirements.

However, this response raises the question as to exactly which reasons a person is being responsive to when she has a merely blameless belief. If her belief that not-P is merely blameless, this means that she has failed to respond adequately to her epistemic reasons and there are excusing conditions surrounding that failure that makes her ineligible for blame. If this is right, then any proposition that is only merely blamelessly acceptable to an agent cannot be associated with an evidence-relative practical reason for her.

This means that when an agent rejects a law on the basis of practical reasons, of which some of the associated propositions are only merely blamelessly accepted, she is not, at least with regard to those merely blamelessly believed propositions, responding to her evidence-relative reasons. At best, she is responding to only some of her evidence-relative reasons. These reasons would be that subset of her practical reasons, the associated propositions of which she can rationally accept, given her evidence. Moreover, insofar as she would not reject the law if she had responded to all her evidence relative reasons, the reasons she failed to respond to are stronger than the reasons she ended up responding to.

To see why, let us reexamine the Immigration case. Suppose that Alice rejected a law that reduced immigration restrictions because she merely blamelessly believed that reducing immigration restrictions would cause locals to lose their jobs. If she had rationally believed, given her evidence, that relaxing immigration restrictions would not cause locals to lose their jobs, she would have supported the law. This can be the case only because the reasons in favour of the law are stronger than the reasons against the law. After all, if the reasons she did not respond to were weaker, then she would have, all things considered, reason to reject the law. But if she had all things considered reason to reject the law, then she would reject the law even if all her beliefs were rationally formed. While such a case is certainly possible, the focus in this section is on laws an agent rejects only because she merely blamelessly forms some belief. In such cases, not only has she failed to respond to at least some of her reasons any reasons she may have responded to are weaker than the reasons she did not respond to.

Moreover, there could be some situations in which all the propositions associated with the putative reasons that an agent cites when she rejects a law are merely blamelessly accepted. In such a case, there are no evidence-relative reasons she can be said to have responded to in rejecting that law. Nevertheless, if blamelessly believing a proposition is sufficient to make that belief reasonable, then since she can blamelessly accept those propositions, she would count as lacking access to the reasons that would justify the law. Hence she would count as reasonably rejecting that law. If there are no reasons that she can be said to have responded to, then epistemic reasonableness does not present any normative demands for her to fulfil.

To elaborate, given that normativity supervenes on reasons, we might say that an agent, Alice, ought to be epistemically reasonable if and only if she has decisive reason to be epistemically reasonable. If being epistemically reasonable amounted to being rational, then we can understand the claim that Alice ought to be rational as saying that she has decisive reason to believe what her reasons require her to believe. The claim that Alice ought to rationally form beliefs is made true by the existence of epistemic reasons available to her. The property of being rational “latches on” to these epistemic reasons. If being epistemically reasonable amounted to being blameless in forming one’s beliefs, then the claim that she ought to blamelessly form beliefs is not made true by anything since blamelessly forming beliefs, as such, does not necessarily involve responding to any reasons. It is true that if Alice does respond correctly to her epistemic reasons, she will, in addition to being rational, also be blameless in her beliefs, but this connection is merely accidental. If Alice has a merely blameless belief, she has not responded properly to her epistemic reasons. Unlike being rational, there are no reasons for the property of being blameless to “latch on” to. Therefore, there cannot be a normative requirement that is satisfied by merely blamelessly forming beliefs. This means that a merely blamelessly formed belief is not a reasonable belief. Therefore, reasonable disagreement has to be rational disagreement between reasonable persons.

The second objection to reasonable disagreement as blameless disagreement is that on such an account, far too few laws can be publicly justified. Enoch presents an objection (2017: 148-149) that nearly every proposition, especially those that would play a role in justifying laws, is subject to blameless disagreement. If so, then for any law, there could be someone who blamelessly rejects the practical reasons that justify the law. If reasonable disagreement is blameless disagreement, then no

law can be justified to everyone because all laws are justified on the basis of doctrines that are subject to reasonable disagreement.

To be clear, Enoch's objection is too strong. While it is true that most propositions are subject to blameless disagreement, it may still be true that some blameless disagreements are not reasonable because they are inconsistent with other non-epistemic requirements of reasonableness: for instance, the commitment to fair terms of cooperation that are mutually acceptable. The claim here is that the acceptance of some substantive normative commitments is constitutive of being reasonable. Therefore, blamelessly rejecting these commitments *ipso facto* makes one unreasonable. If this reply to Enoch's objection is successful, the ubiquity of blameless disagreement does not unacceptably reduce the set of publicly justifiable laws to nil because some epistemically blameless disagreement is still unreasonable.

Nevertheless, we might still preserve a core part of Enoch's objection. Suppose, for the sake of argument, that accepting some core liberal values like freedom and equality is constitutive of being reasonable. These values which are nevertheless fairly abstract are not in and of themselves policy recommendations. Further inferential steps are needed to connect liberal laws and policies to liberal values. It is still possible for people who accept these abstract liberal values to blamelessly, but not rationally, believe that some illiberal and authoritarian laws are just. If this is right, then it may still be difficult to publicly justify many core liberal laws protecting personal liberties.

For instance, the laws that ban same-sex marriage (SSM) are considered by most liberals¹² to be a core instance of illiberal laws. We can take this to mean that morally reasonable persons could not rationally disagree about whether SSM should be legal. However, some morally reasonable persons may, perhaps because they live in very conservative regions and have access only to conservative media, blamelessly believe that SSM should not be legal. They may blamelessly fail to infer that such a ban is not acceptable to other reasonable persons or that it is inconsistent with the reciprocal and equal freedom of everyone. Therefore, laws that recognise SSM would not be justifiable to them and hence not be publicly justifiable.

Similarly, many liberal laws which can be rationally inferred from liberal values may be blamelessly rejected by some reasonable persons. If reasonable disagreement is

¹² See for instance, Lister (2016) for a defence of laws recognising same sex marriage on public reason grounds.

blameless disagreement between reasonable persons, then these liberal laws are not publicly justifiable. This result is significantly less likely if reasonable disagreement was rational disagreement between reasonable persons instead. Therefore, epistemic reasonableness should be rationality and not blamelessness.

The third objection to the account of reasonable disagreement as blameless disagreement between reasonable persons is that on such an account, reasonable disagreement lacks normative significance. Reasonable disagreement is normatively significant in a way that mere disagreement is not. The fact that a given disagreement about a proposition is reasonable makes it inappropriate to use it as a reason to justify laws. Mere disagreement about a proposition, however, does not make justification of laws on the basis of that proposition inappropriate. As such, we might say that reasonable disagreement ought to elicit deference from others about whether to endorse a law.

One commonly cited reason as to why we should defer to others' reasonable rejection of a given law or policy is out of respect for their capacity to respond to reasons (Rawls 2001; Boettcher 2007; Freeman 2007; Larmore 2008; 2015; Gaus 2011; Wall 2016). However, when disagreement is merely blameless, it is possible that at least one agent is not responding to any of her reasons. At the very least, she is not responding to all the epistemic reasons that she epistemically ought to respond to. Therefore deferring to someone who merely blamelessly rejects a law cannot be a way to respect their capacity to respond to reasons. Blameless disagreement, even if between reasonable persons, therefore lacks the kind of normative significance that political liberals think reasonable disagreement should have.

In fact, if reasonable disagreement is normatively significant because we should respect people's capacity to respond to reasons, then reasonable disagreement must entail rational disagreement. The fact that we ought to defer to others' rejection of a law or policy only when their rejection is consistent with them being reasonable suggests firstly that what is being respected is at least in part their capacity to be reasonable and secondly that this is the capacity to appropriately accept or reject laws and policies on the basis of good reasons. If the capacity we are respecting is not the capacity to be reasonable, then there is no reason why we should defer to their rejection when it is or could be reasonable but not otherwise. Likewise, if the capacity we are respecting is not the capacity to appropriately accept or reject laws and policies, then it would be unclear as to why we ought to

defer to someone's rejection of laws and policies, but not defer to them on all their beliefs and projects. It follows from this that reasonableness involves the disposition to respond adequately to one's (evidence-relative) reasons about which laws and policies to accept or reject. It turns out that an agent's evidence-relative reasons are simply those propositions which are supported by her evidence. That is to say, they are propositions that she would be rational in believing, given her evidence. Merely blameless beliefs cannot be associated with evidence-relative reasons; they lack the relevant sort of connection to an agent's evidence. It seems, therefore, that reasonable disagreement must be rational disagreement between reasonable persons.

Summing up, I have presented three objections to the conception of reasonable disagreement as blameless disagreement between reasonable persons. If any of these objections are successful, then reasonable disagreement must be rational disagreement between reasonable persons. It is only if reasonable disagreement implies rational disagreement, is it possible that the standard account of political liberalism presupposes permissivism. The reason for this is that permissivism is a claim about when rational disagreement is possible. If, on the standard account, reasonable disagreement did not entail rational disagreement, then questions about when rational disagreement is possible would be irrelevant to political liberalism.

1.3 Permissivism

As noted earlier, a central concern of the thesis is to show that the standard account of how reasonable disagreement arises presupposes epistemic permissivism, according to which people who share the same total body of evidence can rationally disagree. More specifically,

Permissivism: Given a total body of evidence and proposition possibly more than one doxastic attitude towards the proposition is rationally justified.¹³

By doxastic attitude, I mean something like a belief, disbelief or suspension of judgment. This way of cashing out our doxastic attitudes is particularly coarse grained; it presupposes that there are only three attitudes we might take towards a

¹³ As many authors on this topic have noted, many epistemological theories like van Fraassen's Voluntarist Epistemology, various types of Epistemic Conservatism, Foley/James style instrumentalism, Rawlsian reflective equilibrium, orthodox Bayesianism and coherentism require permissivism to be true (White 2005; Kelly 2010; Ballantyne and Coffman 2011; Schoenfield 2014).

proposition. We might think that this conception of doxastic attitudes does not adequately reflect our actual belief forming practices. We might think that people believe more or less strongly and similarly may strongly, moderately or even weakly disbelieve a proposition. This is a somewhat more fine-grained division of our mental life. Many epistemologists¹⁴ are willing to go even further and represent our doxastic attitudes in terms of precise probabilities. While there is disagreement about the best conception of doxastic attitudes, nothing I argue for in this thesis depends on any specific account being true. However, in order to simplify the presentation of my arguments, I shall, where appropriate, employ the coarse grained account, according to which, there are only three attitudes we might take towards a given proposition. In some parts of the thesis which I shall flag, the discussion will be better served by using a more fine grained conception of doxastic attitudes. In those instances, I shall employ the more fine-grained account instead.

I stipulate that people disagree about a proposition whenever they have different doxastic attitudes towards it. This is most obvious in the case in which Alice believes that P while Betty disbelieves P. For instance, theists, who believe that God exists, certainly disagree with atheists, who disbelieve that God exists. Similarly, it seems that Alice and Betty would disagree if one of them believed that P while the other suspended judgment about P. It is plausible to think that theists and atheists both disagree with agnostics too. If we were to employ a slightly more fine grained account of doxastic attitudes, we can imagine that someone who strongly believed that P would also count as disagreeing with someone who only moderately believed that P. Imagine the following situation:

Bungee Cord: Alice and Betty are both arguing about how safe a given bungee cord is. For instance, suppose Alice is nearly certain it is safe while Betty is only moderately confident. Alice and Betty's argument involves pointing to evidence like the track record of safety for using this type of cord, its physical appearance and specifications.

If they both counted as agreeing about the safety of the cord, their argument would not make sense. However, their argument does not seem pointless. It is potentially momentous, in that Betty might decide not to purchase that cord because she is only moderately confident that it is safe, whereas she may buy it if she is nearly certain that it is.

¹⁴ For instance, Bayesians and all other formal epistemologists.

We might think that with even more fine grained accounts of doxastic attitudes, people who have attitudes that are infinitesimally close to each other would still count as disagreeing. However, this point requires some argument as it seems initially counterintuitive. For instance, Alice may have a confidence that P of 0.75 while Betty has a confidence of 0.76. In such a situation, it is extremely unlikely that such differences in doxastic attitudes could warrant any argument between them. Moreover, normal human agents do not seem equipped to distinguish their doxastic attitudes to such a precise degree. It might therefore seem absurd to suppose that they disagree.

The claim that Alice and Betty agree depends on the intuition that they cannot disagree if their doxastic attitudes are so similar as to be indistinguishable from each other. However, indiscriminability of doxastic attitudes does not entail agreement. To see why, we merely need to consider a few more cases. Suppose another agent, Celia, has a confidence of 0.77 that P , Darla has a confidence of 0.78 and so on all the way to Zelda who has a confidence of 1 that P . Notice that indiscriminability is not transitive. Alice's attitude is indiscriminable from Betty's and Betty's is indiscriminable from Celia's and so on all the way to Zelda. At some point Alice's attitude is discriminable from a third person's attitude even if both Alice and the third person's attitudes are indiscriminable from someone whose attitude is in between the two. At the very least, Zelda's attitude is obviously discriminable from Alice's. After all, being perfectly confident that P is very different from supposing that there is a 1 in 4 chance that P is false. However agreement is transitive. If Alice agrees with Betty, Betty with Celia, Celia with Darla all the way to Zelda, then Alice must agree with Zelda. However, if Alice has a confidence of 0.75 and Zelda a confidence of 1, then they do not agree. Therefore, there has to be at least one instance in the chain whereby one agent disagrees with another agent even though their doxastic attitudes are indistinguishable from each other. However, if we think that, for instance, Mary who has a confidence of 0.87 disagrees with Nina who has a confidence of 0.88, then consistency requires us to also think that Alice disagrees with Betty, Betty with Celia and so on. After all, there is no reason to think that Alice agrees with Betty but Mary does not agree with Nina when Alice's attitude is as close to Betty's as Mary's is to Nina's. Therefore, Alice cannot agree with Betty no matter how arbitrarily close to each other their doxastic attitudes are. They agree if and only if they have identical doxastic attitudes.

With this account of disagreement, we can note that permissivism does not claim that rational disagreement is possible on any question no matter how complete your

evidence is. The claim is more modest, namely, that there are at least some cases in which rational disagreement is possible given the same evidence. This is because if permissivism were true, there could be cases in which two people who have the same evidence could have different doxastic attitudes towards a proposition and both would still be rational. If permissivism were false, then if two people who have the same evidence disagree, then at least one of them must be irrational.

1.4 Chapter Summaries

One of the central aims of this thesis is to demonstrate that the standard account of political liberalism presupposes permissivism. I shall do this by showing that the burdens of judgment, which are invoked to explain reasonable disagreement in the standard account, presuppose permissivism. In the first part of this argument I shall explain what it means to have and share evidence; in the second, I shall explain for each of the listed burdens of judgment, how each burden presupposes that people can reasonably disagree even when they have the same evidence. These arguments will be presented in the Chapter 2.

In the rest of the thesis, I shall argue that permissivism is false by showing that the uniqueness thesis, the negation of permissivism, is indeed true.

Uniqueness: Given a total body of evidence and proposition, at most one doxastic attitude towards that proposition is the rationally justified one.

In Chapter 3, I shall initially distinguish between intrapersonal and interpersonal uniqueness. Intrapersonal uniqueness is the claim that given the evidence, at most one doxastic attitude towards a proposition would be rational for a given agent. This would be compatible with permissivism because even if intrapersonal uniqueness were true, different attitudes could be uniquely justified for different agents even if they all share the same total body of evidence. Interpersonal uniqueness would make the stronger claim in that the doxastic attitude which is justified, given the evidence, would be the same for all agents.

I shall then argue that any plausible account of permissivism will be consistent with intrapersonal uniqueness even if not with interpersonal uniqueness. This will involve arguing that any theory which violates intrapersonal uniqueness is open to Roger White's (2005; 2014) Arbitrariness Objection and his Evidence Pointing Problem

which I shall present. According to the Arbitrariness Objection, if my evidence were permissive and made more than one doxastic attitude permissible for me, then choosing an attitude by flipping a coin would be just as good as reasoning on the basis of my evidence. However, since flipping a coin is an arbitrary and irrational way to form beliefs, so also must reasoning on the basis of permissive evidence. According to the Evidence Pointing Problem, the metaphysics of evidential support preclude any body of evidence from, at the same time, supporting a proposition to different degrees. That is to say, evidential support is unidirectional. If evidential support were unidirectional, then uniqueness (at least in its intrapersonal version) must be true.

I shall subsequently present and address two objections to intrapersonal uniqueness. The first is Meacham's objection to the arbitrariness objection according to which arbitrariness as such is, contra White, independent of questions of uniqueness or permissiveness. The second consists of a series of counterexamples in which people can justifiably have any doxastic attitude towards a proposition because those propositions are self-fulfilling.

After defending intrapersonal uniqueness, I shall present what I take to be the most plausible version of permissivism. On Schoenfield's version of permissivism which is consistent with intrapersonal but not inter-personal uniqueness, what is epistemically rational for a person to believe is not only dependent on her evidence, but also on what epistemic standards she holds. This claim is not only consistent with the second burden of judgment, it goes a bit further: People can rationally disagree about a proposition given the same total body of evidence if and only if they can permissibly disagree about the weight of various evidential considerations. This makes the question of whether permissivism is true depend on whether people can permissibly have different epistemic standards. I shall describe two variations of this how this might be possible. On the first variation, agents can permissibly choose between multiple epistemic standards. I call this view intrapersonal permissivism about epistemic standards. On the second variation, there is one epistemic standard that is suited for each agent but different standards may be best suited for different agents. I call this view intrapersonal uniqueness about epistemic standards.

In the remaining chapters, I shall argue that given that intrapersonal uniqueness is true, so is interpersonal uniqueness. In the Chapter 4, I shall argue that White's Arbitrariness Objection and the Evidence Pointing Problem can be extended against

intrapersonal permissivism about epistemic standards. In the first of two parts of this chapter, I shall rehearse Bob Simpson's extension of the Arbitrariness Objection to epistemic standards. The argument, roughly, is that if, for an agent more than one epistemic standard would have been rational for her to adopt, then reasoning from the evidence on the basis of her own epistemic standard is still arbitrary. In the second part of the chapter, I shall argue that the same metaphysical considerations that make evidential support unidirectional also make it impossible that more than one epistemic standard is rationally permissible for an agent. If the two arguments in this chapter succeed, we can conclude that people can permissibly disagree about the strength of various evidential considerations only if some other burden of judgment is in play. This leaves the second variation of permissivism which is consistent with intrapersonal uniqueness about standards.

In each of the next three chapters, I address one version of permissivism consistent with intrapersonal uniqueness about standards. That is to say, I shall present and object to versions of permissivism according to which the standard which is most suited for any given agent depends on some feature of the agent, like her cognitive capacities, her epistemic values or her prior credences.

In the Chapter 5, I shall discuss the view according to which the epistemic standard that is appropriate for an agent depends on her cognitive capacities. I shall present Simpson's exposition of this view and show that it ultimately depends on the claim that epistemic standards can be lowered if one's cognitive capacities are limited. I then show how this claim is implicit in Rawls's first, third and fifth burdens of judgment and is present more explicitly in Gaus's account of rational disagreement. I shall then argue that any conception of rationality in which epistemic standards can be lowered in this way has more implausible implications than the alternative. If my argument is successful, epistemic standards cannot be lowered because people have limited cognitive resources. It follows that epistemic standards cannot vary on the basis of differing cognitive capacities.

In Chapter 6, I shall discuss a version of permissivism according to which the epistemic standard which is most appropriate for an agent depends on how she values the twin epistemic goals of acquiring truths and avoiding falsehoods. This version of permissivism aligns with one way of interpreting the fourth burden of judgment. The thought here is that as per the fourth burden, one of the ways in which people's different backgrounds can cause them to rationally disagree is by causing them to have different attitudes to epistemic risk and different ways of

valuing the twin epistemic goals. I shall argue that on any such version of epistemic instrumentalism, all except a few ways of valuing the two goals, result in committing the agent to two different likelihoods about the proposition in question and are hence, impermissible. Moreover, all of the permissible valuations generate the same epistemic standard. Therefore, even if there are different permissible ways of valuing these goals, they do not license different epistemic standards.

In Chapter 7, I shall discuss a version of permissivism according to which the epistemic standard that an agent ought to adopt depends on her pre-existing beliefs. This aligns with the second way in which the fourth burden can be a source of rational disagreement: People's different backgrounds might cause people to rationally disagree by causing them to have different initial beliefs. This version relies on some sort of epistemic conservatism. I shall present what I take to be the standard version of epistemic conservatism and present some considerations in its favour and show why it is ultimately implausible. I shall then discuss attempts to weaken Standard Conservatism, and show that on any version of conservatism that is weak enough to avoid the costs of Standard Conservatism, no permissive case can be constructed. In addressing these three versions of permissivism, I shall have demonstrated that there is no plausible epistemological theory according to which people ought to adopt different epistemic standards. If this is resolved, then permissivism is false.

In Chapter 8, the concluding chapter, I shall summarise the main arguments in the thesis, flag certain unresolved issues and propose avenues for further research.

Chapter 2: The Burdens of Judgment and Permissivism

Background

Permissivism is the claim that more than one doxastic attitude about a proposition can be rationally permissible given a particular total body of evidence. I will take it to be the case that people disagree whenever they have different doxastic attitudes towards the same proposition. Uniqueness, or the negation of permissivism, as such, does not imply that people cannot rationally disagree, only that people who have the same total body of evidence cannot rationally disagree. As alluded to earlier, the standard version of political liberalism attributes the existence of rational disagreement among morally reasonable persons to the burdens of judgment. My aim in this chapter is to show that the burdens of judgment, as standardly specified, imply permissivism. Insofar as the burdens of judgment or anything like them presuppose permissivism, if my argument in subsequent chapters against permissivism is successful, then the standard version of political liberalism is not tenable.

In order to show that the burdens of judgment presuppose permissivism, I intend to show that the burdens are most plausibly interpreted as implying that rational disagreement can arise from sources other than the possession of different evidence. In approaching this, I shall first explain what I mean by people sharing or having the same evidence, and then, for almost all of the listed burdens, show that each burden implies that rational disagreement can arise even when people share the same evidence with respect to a proposition. First, I shall explain what I mean by sharing evidence.

Permissivism is a principle that has been traditionally invoked in peer disagreement debates: When two epistemic peers who disagree encounter each other, can they continue rationally disagreeing or should their beliefs converge? The conciliationist view, according to which disagreeing peers' doxastic attitudes ought to converge, has been defended on the grounds that uniqueness is true. If at most one doxastic attitude is justified given the evidence, then it cannot be rational for both parties to maintain their beliefs in the face of disagreement (van Inwagen 2010; Kelly 2010). If, on the other hand, permissivism is true, this opens up the possibility that disagreeing epistemic peers need not converge in their attitudes. Uniqueness and permissivism become relevant to the peer disagreement debate only if epistemic peers are defined as sharing a total body of evidence and being equally competent at evaluating it (Lackey 2010; Kelly 2010; Matheson 2018). While many authors in

the peer disagreement literature do define epistemic peers as sharing a total body of evidence, little has been said about what this means.

2.1 Evidence

By the term evidence, I mean anything that could bear, one way or another, on the proposition at hand and which is accessible to the agent. This is compatible with different views about what counts as evidence. On some views¹⁵, evidence is propositional, whereas on other views, subjective experiences lacking propositional content can also count as evidence (Conee and Feldman 2008). For ease of presentation, most of the cases I present will involve propositional evidence, but, that is not to be taken to mean that I endorse any particular conception of evidence as being necessarily either propositional or non-propositional. I do, however, take evidence to be mental (Ibid) and carry some positive epistemic status. Thus, material I have not read would not be counted as evidence if I have not read that report. By positive epistemic status, I mean that not just any belief or mental state counts as evidence. It may be that only knowledge (Williamson 2000), or justified beliefs (Conee and Feldman 2008), or beliefs formed through a reliable process (Comesana 2010; Goldman 2011) count as evidence. While nothing I discuss in this thesis hangs on which positive epistemic status makes any given mental state count as evidence, I shall, unless otherwise stated, treat E as evidence for an agent if she knows that E. This usage in this thesis is not meant to commit myself to the claim that only factive mental states are evidence; only that some factive mental states, perhaps in addition to non-factive mental states, do count as evidence.

A key clarification of the “bears on” relation is appropriate here. An agent may be said to possess a body of evidence E even if, in a manner of speaking, E is not evidence for her. On some elocutions, E is evidence about a proposition P for an agent only if coming to know that E could in some possible situations give *her* a reason to revise her doxastic attitude towards P. However, I shall define the “bears on” relation more loosely: E bears on P if and only if coming to know that E could in some situation give *someone* a reason to revise her doxastic attitude towards P. Therefore, to say that an agent, Alice possesses a body of evidence E regarding P is to say that she knows that E and that coming to know that E could in some situations give someone, not necessarily Alice, reason to revise her doxastic attitude towards P. I shall explain this with the following example:

¹⁵ See Timothy Williamson (2000: 184).

Biblical Geology: John and Mary both know that a literal reading of the Bible dates the age of the earth as being less than 10 000 years¹⁶. For illustrative purposes, let us suppose that we can assign a number between -1 and 1 to the strength of a consideration. When the strength of a consideration is less than 0, it counts against the proposition and when it is more than 0, it counts in favour. Suppose that John thinks that the strength of biblical considerations is 0.1 while Mary thinks that it is 0. This eventually leads them to take different doxastic attitudes towards the proposition that the earth is more than 10 000 years old.

Assuming for the sake of argument that John's weightage of biblical considerations is rationally permissible, then, in a certain manner of speaking, the Bible is evidence about the age of the earth for John. By that same measure, it would seem that the Bible is not evidence about the age of the earth for Mary. However, this sense of being evidence need not be the most appropriate for our purposes. According to my definition, the Bible bears on the age of the earth if and only if there is someone, in this case John, for whom there is some situation where learning what the Bible says gives him some reason to believe that the earth is less than 10 000 years old. Correspondingly, Mary knows that the Bible claims that the earth is less than 10 000 years and as per assumption, what the Bible says bears on the proposition in question.

This results in the counter-intuitive claim that the Bible is part of Mary's body of evidence that pertains to the age of the earth even when she justifiably takes it to be irrelevant. On the other hand, given our concern for whether evidence is shared, this way of cashing out evidence is compatible with intuitive claims about when two people share a body of evidence. For instance, to claim that John and Mary do not share evidence implies that one of them does not know or is not in a position to know what the Bible says about the age of the earth. This is clearly false. They both know what the Bible says; it is merely that one of them does not think what the Bible says is relevant. Therefore, it seems as if we must conclude that the Bible is part of Mary's evidence even if there is no situation in which knowing what the Bible says gives her a reason, one way or the other, to revise her beliefs about the age of the earth. Moving forward, in discussing what it means for two people share a body of

¹⁶ The claim here is only that they know that the Bible says that it is less than 10 000 years old, not that they are justified in believing that it is less than 10 000 years old on the basis of the Bible.

evidence, I shall implicitly presuppose that evidence regarding a proposition is a mental state possessed by the person and which, in some possible situation, could give someone a reason to revise their doxastic attitude towards that proposition.

As mentioned earlier, one key concern, in this chapter, is to account for what it means for evidence to be shared, or for two people to have the same piece of evidence. As such, evidence must, in principle, be shareable and, where questions of uniqueness or permissivism are at stake, shared. However, if evidence refers to some kind of mental state, what does it mean to share a mental state? As an initial definition, I shall propose that two people, John and Mary, share a body of evidence E if and only if John has a body of evidence which is sufficiently similar to a body of evidence possessed by Mary. By sufficiently similar, I mean the following:

Sufficiently Similar: One person's, John's, body of evidence E1 is sufficiently similar to another person's, Mary's, body of evidence E2 with regards to a proposition P if and only if a) the set of doxastic attitudes towards P it would be rational for John to adopt, given E1, is identical to the set of attitudes towards P it would be rational for him to adopt, given E2 instead and b) the set of attitudes towards P it would be rational for Mary to adopt, given E2, is identical to the set of attitudes towards P it would be rational for her to adopt, given E1 instead.

There are two things we may notice about the above definition. The first is that sufficiently similar bodies of evidence are not identical. The second notable feature is that bodies of evidence are individuated according to which doxastic attitudes they rationalise. I shall explain and defend each feature in turn. Let us first consider the non-identity of similar evidence. Notice that given Sufficiently Similar, even if E1 and E2 are sufficiently similar, if John were to acquire E2 in addition to E1, it is possible that he may learn something new and thereby acquire additional reasons to believe one way or the other about P. Consider the following case:

Raven: John observes that 50 ravens are black. Mary observes that an entirely different set of 50 ravens are black. John forms a high degree of confidence in the proposition that all ravens are black on the basis of his observation. Mary forms a slightly lower degree of confidence in the proposition that all ravens are black.

John's and Mary's bodies of evidence are sufficiently similar to each other. If John had Mary's evidence instead, he would be just as rational in forming the same doxastic attitude he currently holds in response to his current evidence. The reason for this is that either way, he has observed 50 black ravens. However, if he had acquired Mary's evidence in addition to his own, he would have an even higher degree of confidence that all ravens are black because he would have now observed 100 black ravens. Nevertheless, even though each agent's body of evidence is distinct, the two are sufficiently similar for each to be considered as sharing the body of evidence with the other.

Our goal in attempting to provide an account of shared evidence is to account for how rational disagreement is possible. No explanation as to how John and Mary can rationally disagree can appeal to the fact of different bodies of evidence. After all, both have only observed 50 black ravens. There is no reason why the fact that the ravens they observed were different should make a difference to which attitudes are rationalised by the evidence. It follows then that any explanation for how they could rationally disagree is going to be identical to the explanation that would be called for if they had observed the same set of ravens. Therefore, we should treat two people as sharing evidence even when their evidence is not identical.

As mentioned earlier, the second notable feature of Sufficiently Similar is that bodies of evidence are individuated according to which attitudes they rationalise. Two bodies of evidence are sufficiently similar as long as they rationalise the same distribution of doxastic attitudes over each agent possessing said evidence. Little has been said about how to individuate bodies of evidence. Where evidence is propositional, this lack of discussion is understandable as a reductive account of shared evidence is readily available. John and Mary share the same body of evidence E if and only if they both know that E or justifiably or reliably believe that E. Thus, there seems to be a clear sense in which propositional evidence can be shared. Things are less clear cut for non-propositional evidence like perceptual experience.

I shall present two reductive accounts of sharing non-propositional evidence and show that neither account captures all our intuitions about when we share a body of evidence. The first account is that people share a body of perceptual evidence if and only if they have identical perceptual experiences. On the second account, people share a body of evidence if and only if they have perceptual experiences of the same object. I shall address each account in turn.

The claim that people share evidence if and only if they have identical perceptual experiences is implausible because it is under-inclusive. We can easily imagine situations in which people have non-identical perceptual experiences, but have the same body of evidence. Consider the following case:

Stopwatch: John and Mary are measuring the period of a pendulum with a stopwatch. John is standing to the left of Mary and they are both looking at the stopwatch that she is holding. As such, the stopwatch is slightly to the right of the centre of John's visual field, while it is left of centre in Mary's visual field.

In Stopwatch, both John and Mary can clearly see the reading on the stopwatch and therefore have the same evidence about the period of the pendulum. Considering that people can have the same body of evidence when their perceptual experiences are slightly different, we might attempt to weaken our definition by claiming that people share a body of (non-propositional) evidence if and only if their perceptual experiences are highly similar to each other. However, two highly similar perceptual experiences might count as the same body of evidence with respect to one proposition, but not another. Consider the following case:

Chipped Workbench: The workbench that John and Mary are using for their experiment is chipped. Coincidentally, Mary is holding the stopwatch in such a way that the watch blocks John's view of the chip on the workbench but not Mary's. Therefore only Mary can see that the workbench is chipped.

In Chipped Workbench, the two visual experiences which counted as the same evidence about the period of the pendulum count as different bodies of evidence when the proposition is about whether the workbench is chipped. This suggests that there is no fixed degree of similarity between two perceptual experiences which suffices to make them count as the same body of evidence. If so, then there is no general explication of non-propositional evidence in terms of the content of perceptual experience.

According to the second reductive account, people have the same evidence if and only if they have a perceptual experience of the same object. As with the first reductive account, we can imagine cases where two people have perceptual experiences of the same object, but they clearly have different bodies of evidence. Consider the following case:

Cat/Jacket: There is an object in the garden which both John and Mary are looking at. John and Mary are standing some distance from each other. Due to the way the light hits the object relative to John's position, it looks like a cat to John. Due to the way the light hits the object relative to Mary's position, it looks like a fur jacket to her. If Mary were standing in John's position, the object would look like a cat to her. Likewise if John were standing in Mary's position, it would look like a jacket to him.

In Cat/Jacket, we see that even though John and Mary both have a visual experience of the object, they have different evidence. After all, there is no other way to explain why the object would have looked like a cat instead of a jacket to Mary if she had been standing where John was.

Summing up, I have presented two reductive definitions of shared evidence and shown that both accounts fail to account for some of our intuitions about when we share evidence. To be clear, I am not claiming that there is no metaphysical reduction of the notion of evidence to perceptual experience, only that no reductive definition of shared evidence is easily available. If this is right, then, given that Sufficiently Similar adequately captures our intuitions about when two people share a body of evidence, we are now in a position to specify when people share a total body of evidence:

Total Evidence Sharing: Two people, Mary and John, share a total body of evidence regarding a proposition P if and only if all the evidence John has regarding P is sufficiently similar to all the evidence regarding P that Mary has.

Here, Total Evidence Sharing relies on Sufficiently Similar to define what it means when the piece of evidence one person has is the same as a piece of evidence another person has.

Permissivism is the claim that given a total body of evidence, more than one doxastic attitude towards a proposition can be justified. That is to say, two people who share a total body of evidence can rationally disagree about a proposition. This is ensuing from the thought that people can rationally disagree about a proposition only if more than one doxastic attitude towards it is justified. On the standard account of political liberalism, reasonable disagreement is attributed to the burdens of judgment. They are causes of disagreement which do not impugn the

reasonableness of the disagreeing parties. The following six causes are cited by Rawls (1993), namely: that the evidence is difficult to assess, that people disagree about the weight of various considerations, that the concepts we use are subject to vagueness, that we have different background experiences, that moral questions are complicated with considerations on both sides of the issue and that people function in a limited social space. In the previous chapter, I argued that reasonable disagreement was rational disagreement between reasonable persons. As such, showing that the standard account of political liberalism presupposes permissivism requires, first, providing an account of what it means to share a total body of evidence and, secondly, showing how most of the burdens of judgment specify conditions for rational disagreement that are compatible with the disagreeing parties sharing the same total body of evidence with regards to the proposition in question. Having clarified what it means for two people to share a total body of evidence, I shall now proceed to examine each of the listed burdens, in turn, in order to show that at least four of the six burdens attribute rational disagreement to some factor other than having different evidence. If rational disagreement is attributable to something other than having different evidence, then when those burdens affect the judgment of disagreeing parties, people can rationally disagree even if they had the same total body of evidence.

2.2 The First Burden

Consider the first burden of judgment:

“a. The evidence – empirical and scientific – bearing on the case is conflicting and complex and thus hard to assess and evaluate.” (ibid: 56)

The implicit claim of the first burden is that if the evidence bearing on a case is hard to assess and evaluate, agents who have that evidence may permissibly disagree, that is, take different doxastic attitudes towards the same proposition.

To elaborate, consider a body of evidence consisting of a hundred propositions, P1 to P100. Suppose that two agents, Alice’s and Betty’s total body of evidence consists of P1 to P100 because they each know all these propositions. Suppose, further, that some of these propositions represent polling data from surveys which suggest that candidate A will win the upcoming elections. Other propositions represent polling data from surveys which suggest that candidate B will win instead. Given the large number of propositions containing conflicting information, the

evidence can be considered conflicting and complex. According to the first burden, given the complexity of the evidence, there may not be any single correct way of putting together all this conflicting information. If so, then Alice could permissibly interpret her evidence as supporting the proposition that candidate A will win while Betty could permissibly interpret that same body of evidence as supporting the proposition that candidate B will win.

It might be the case that in virtue of Alice and Betty being able to permissibly disagree given that they know P1 to P100, they may come to know different propositions. On some accounts of evidence, this might result in them having different bodies of evidence. For instance, if an agent's evidence consists of propositions that she knows, since Alice and Betty now know different propositions, their total body of evidence differs¹⁷. However, we should not take this to mean that the evidence being difficult to assess is incompatible with people sharing a total body of evidence. At the time before they have made different inferences, Alice and Betty do have the same total body of evidence.

Clearly then, according to the first burden, the difficulty of evaluating evidence could licence rational disagreement even when the disagreeing parties share the same total body of evidence. It follows that the claim made by the first burden is that permissivism is true, at least in part, because evidence can sometimes be difficult to assess.

To see why, suppose that permissivism were false. Then consider again Alice and Betty's situation in which the evidence regarding who will win the election was conflicting and complex and hard to evaluate. If permissivism was false, then it would not be rational for them to disagree. After all, uniqueness just is the claim that at most one doxastic attitude towards a proposition is justified given a total body of evidence. If the difficulty of evaluating the evidence makes it rational to disagree, then permissivism has to be true.

As I explain at greater length in Chapter 5, the difficulty of evaluating the evidence can make rational disagreement possible only if whether a belief counts as rational, given the evidence depends on the agent's cognitive capacities. Kornblith (1983)

¹⁷ As Conee and Feldman (2008) note, if we count the propositions that we infer from our other evidence as evidence too, we end up double-counting what reasons we have to believe our target proposition. For instance, if my evidence consists of P and P entails Q which in turn entails R, inferring that Q from P does not give me an additional reason to believe R. This is a reason to suppose that propositions which we infer from our existing evidence are not themselves evidence.

and Rosa (2012; 2016) argue that we are permitted to not make otherwise obligatory inferences if we are too cognitively limited to make those inferences. People could rationally disagree given the same body of evidence if they were cognitively limited to different degrees. Only if some account like this is successful can the difficulty in assessing the evidence make disagreement rational.

2.3 The Second Burden

The second burden of judgment is more explicit about presupposing that reasonably disagreeing parties can share the same evidence.

“b. Even where we agree fully about the kinds of considerations that are relevant, we may disagree about their weight, and so arrive at different judgments.” (Rawls 1993: 56)

The second burden concerns situations in which disagreeing parties agree what the relevant considerations are, that is, situations in which disagreeing parties share the same total body of evidence with regards to a proposition and agree that these considerations matter. The considerations for and against a proposition, insofar as they are accessible to an agent, just are her evidence for and against that proposition. To see why, we can first note that plausibly any given piece of evidence regarding a proposition is a consideration for or against that proposition. It is unclear how something could be a relevant consideration, but not evidence. It follows that if people agree what the relevant considerations are, they share the same evidence.

To see how this can be the case, consider the situation where they do not share the same evidence. Suppose two people, Mary and John have different bodies of evidence for a proposition P. Suppose John has evidence for P that Mary does not have. Then, this would be an additional consideration that John thinks bears on P but which Mary does not. This makes clear that any difference in our evidence creates differences in what we may permissibly take to be the relevant considerations. Therefore, if John and Mary permissibly agree on what the relevant considerations are, they have to have the same evidence. The second burden of judgment therefore implies that even in cases where people share the same total body of evidence, they could rationally disagree because they permissibly disagree about the strength of certain parts of their evidence. This certainly presupposes permissivism.

One question that might be asked is how two people who have the same total body of evidence could come to disagree about the strength of a given consideration. After all, it could be the case that even though two people share a given piece of evidence even when they disagree about how strongly it supports a proposition, they could disagree about the strength of the evidence only because their total body of evidence differs in some other respects. For instance, if John and Mary disagree about how important Biblical evidence is with regards to geological questions, they may turn out to share the Biblical evidence, but differ with regards to evidence about how reliable the Bible is.

However, this need not be the case. On Schoenfield's view (2014) people can permissibly disagree about the strength of their evidence because they have different epistemic standards. Epistemic standards consist of epistemic norms, which are rules and principles regarding when to form certain doxastic attitudes. On some epistemological accounts (Goldman 2010), epistemic norms are source authorisations which instruct the agent to treat some sources of information as reliable. For instance, John might accept a norm that tells him to treat the Bible as a reliable source of information, while Mary accepts a norm that tells her to treat the Bible as fictional. Other epistemic norms might say that the BBC, but not Fox News, is a reliable news source, or vice versa. On other accounts, they might be more general rules like "if it visually seems to you that P, believe that P" or "do not trust your intuitions about when someone is being sincere". In general, epistemic norms will be rules of the form "if your evidence is E, form attitude A towards proposition P". Any plausible account of permissivism will thereby provide some explanation as to how people can have different epistemic standards without having different evidence about which of the constituent norms are correct. In Chapter 3, I shall explain in greater detail how the claim that agents can have different epistemic standards addresses two powerful objections to permissivism: The Arbitrariness Objection and The Evidence Pointing Problem. If this is right, then the second burden of judgment represents not just one more specific cause of disagreement, but provides a framework by which other burdens of judgment generate rational disagreement. For instance, if, as per the first burden, limitations to cognitive capacities can make rational disagreement possible, it can do so only by making it permissible for people with different cognitive capacities to have different epistemic standards. Addressing whether the second burden of judgment can in fact cause rational disagreement would involve addressing whether any of the other burdens of judgment permit people to have different epistemic standards.

2.4 The Third Burden

I shall now turn to the third burden according to which:

“c. To some extent, all our concepts, and not only moral and political concepts, are vague and subject to hard cases; and this indeterminacy means that we must rely on judgment and interpretation (and on judgments about interpretations) within some range (not sharply specifiable) where reasonable persons may differ.” (Rawls 1993: 56)

The third burden of judgment claims that people could come to reasonably disagree because the concepts that apply in a given case are vague and indeterminate. The thought here is that many of the predicates we apply in our ordinary language like old, tall or red are vague. For instance, Jeanne Calment, the longest lived person in recorded history, died at the age of 122 years¹⁸. When she died she was old. Presumably she was not always old. There was some time in the past when she was not old, and some time near the end of her life that she was. It is unclear, however, when she stopped being not old and became old. If she was old on her 70th birthday, was she old the day before? What about the day before that? If we iteratively ask the question, we eventually reach a point when she was clearly very young. The concept of old, then, seems to have fuzzy edges that make it difficult to determine whether the predicate “is old” applies in certain borderline cases. Moreover, while there clearly are borderline cases, vagueness is iterative and the boundary between clear-cut cases and borderline cases is vague. Rawls’s claim with regards to the third burden is that in borderline cases, people can rationally disagree about whether a person at a given age is old because of the vagueness of the predicate, “is old”.

One plausible account¹⁹ of vagueness is epistemic. On Williamson’s account (1994), the predication of every borderline case is either true or false. Vagueness arises because our capacities to determine whether a given predicate does apply to

¹⁸ Williamson (1994) invokes the example of Rembrandt. However, Rembrandt died at the age of 63, and 63 does not necessarily seem old. This proves the point that the predicate “is old” is vague.

¹⁹ There are other accounts of vagueness like Subvaluationism (Hyde and Colyvan 2008; Corberos 2011), according to which, borderline cases are both true and false; and Supervaluationism (Fine 1975; Lewis 1982), according to which borderline cases are neither true nor false. The fact that moral questions tend to be difficult makes epistemicism a better explanation for vagueness when it comes to moral questions than for questions like whether someone is old or when something is a heap.

a case are limited. Williamson leaves open exactly which capacities are involved in failing to know whether a given predication of an object is true. It might, for instance, be that there is some property of oldness or thinness which we are able to perceive in clear-cut cases but not in borderline cases even when the property is present. Alternatively it might be more akin to our inability to determine if Goldbach's conjecture, according to which every even number greater than 2 is the sum of two prime numbers, is true (ibid: 209-212). If the latter, then vagueness is, like with the first burden, a result of the evidence being too complicated to assess. As with the discussion of the first burden, this complexity is compatible with people sharing the same total body of evidence. On this interpretation, the claim made by the third burden is that even if everyone had the same evidence, the difficulty in interpreting the complex evidence generates rational disagreement about when a predicate applies to a borderline case.

Even if the cognitive limitation is a matter of being unable to perceive oldness or thinness in borderline cases, it does not follow that people have different evidence about whether the predicate applies. As Williamson notes, vagueness exists even within a community of competent language users. Given that, according to Williamson, competent language users have a shared corpus of knowledge about when to use the word thin or old (ibid: 211), we have good reason to think that a community of competent language users share a total body of evidence regarding the meaning of terms like thin or old as used in the community. The explanation for rational disagreement, then, is that this total body of evidence underdetermines the conclusion.

One thought that might help explain how the evidence can underdetermine the conclusion is that limitations to our perceptual capacities mean that we can, in some circumstances, fail to perceive properties or only partially perceive them even when they are instantiated. Given that perceptual evidence is non-propositional, the information conveyed can itself contain a degree of imprecision because the property is only partially perceived. The thought, then, is that insofar as the evidence itself contains some degree of imprecision, there are some permitted precisifications of this evidence. Any conclusion that can be permissibly drawn from one of the permitted precisifications of the evidence is a rationally permissible inference from the evidence. Thus, on this interpretation of the third burden, even if two people share the same total body of evidence, as long as there is some degree of imprecision in the evidence itself, they can rationally disagree about the conclusion. Summing up, given an epistemic account of vagueness, rational

disagreement that pertains to conceptual vagueness is consistent with the disagreeing parties having the same total body of evidence.

2.5 The Fourth Burden

According to the fourth burden of judgment:

“d. To some extent (how great we cannot tell) the way we assess evidence and weigh moral and political values is shaped by our total experience, our whole life up to now; and our total experiences must always differ. Thus, in a modern society with its numerous offices and positions, its various divisions of labor, its many social groups and their ethnic variety, citizens’ total experiences are disparate enough for their judgments to diverge, at least to some degree, on many if not most cases of any significant complexity.” (Rawls 1993: 56-57)

The fourth burden claims that people can rationally disagree because they have different backgrounds. One way to understand the effect of different backgrounds is to suppose that due to different life experiences, people acquire different evidence. This evidence might be object-level evidence. For instance, a climate scientist is more likely to encounter meteorological data than the lay-person. The evidence might also be higher-order evidence about whether an epistemic norm is correct. In the Biblical Geology case, someone like John who comes from a very conservative religious background might receive testimonial evidence that the Bible is a reliable source of information while someone like Mary, who is raised in a more secular environment or come from a different religious background will not receive such testimonial evidence. This might make it rational for people from religious backgrounds to attach greater weight to what the Bible says²⁰. Even if permissivism is false, people can still rationally disagree if they have different evidence.

However, causing people to acquire different evidence is not the only way in which having different backgrounds might bring it about that people rationally disagree. For instance, having different backgrounds might cause agents to value acquiring truths and avoiding falsehoods differently. Some people’s moral beliefs depend on certain religious claims. For instance, a conservative’s opposition to abortion in

²⁰ See, for instance, Goldman (2010) who argues that people might reasonably disagree if they get different evidence about which epistemic norms are correct. This type of reasonable disagreement is not the focus of this thesis because it involves people having different evidence.

nearly all circumstances might stem from the belief that God ensouls fetuses at the moment of conception. Two agents might disagree about the existence of God because they have been raised to have different attitudes towards epistemic risk. For instance, someone raised in a more sceptical background may require more evidence before they are prepared to believe a proposition. On the other hand, someone who has been raised in a way that inculcates credulity would correspondingly discount the consequences of having a false belief. These different ways of valuing various epistemic goals might license different epistemic standards and hence different doxastic attitudes even when the total body of evidence is the same. I shall address this possibility in Chapter 6.

In addition, having different backgrounds may also cause people to have different prior doxastic attitudes towards the proposition in question. Along a similar vein of thought, Gaus (2011) suggests that the order in which an agent encounters evidence can affect which attitude towards a proposition she is justified in adopting. To illustrate, suppose that a body of evidence, *E*, can be split into two components *E*₁, which supports the proposition *P*, and *E*₂, which supports not-*P*. Gaus supposes that there might be a case where Alice encounters *E*₁ first and comes to believe that *P* and then encounters *E*₂. However, *E*₂ is insufficient to shift her belief that *P*. By contrast, Betty encounters *E*₂ first and believes that not-*P* and subsequently encounters *E*₁ but *E*₁ is insufficient to shift her belief that not-*P*. Gaus's claim ultimately relies on a version of epistemic conservatism: Less evidence is needed to justify belief in a proposition if the agent already believes it than if the agent does not believe it. Therefore, according to the conservative, the same total body of evidence could be sufficient to justify belief in a proposition if an agent already believes the proposition but insufficient if she does not. I shall address this along with other versions of permissivism based on conservatism in Chapter 7. In showing that conservatism could not ground any version of permissivism, I would have shown that people's different backgrounds cannot give rise to rational disagreement by virtue of having merely started off with different beliefs even if have the same total body of evidence.

2.6 The Fifth Burden

According to the fifth burden of judgment:

“e. Often there are different kinds of normative considerations of different force on both sides of an issue and it is difficult to make an overall assessment.” (ibid: 57)

The fifth burden of judgment is analogous to the first burden; except where the first burden is about the difficulty of judging the normative epistemic force of all one's evidential considerations, the fifth burden is about the difficulty of judging the normative moral force of all the moral considerations. Recall that according to the first burden, the complexity of scientific and empirical evidence makes it difficult to assess the overall strength of one's evidence. This difficulty licenses disagreement about the target proposition as rational. Similarly, according to the fifth burden, rational disagreement about a moral question is made possible because of the greater complexity of different moral considerations. To see the distinction more clearly, I shall present an instance of each type of question.

An instance of the first burden in action is when trying to determine whether increasing the minimum wage exacerbates unemployment. The evidence regarding this is complex and conflicting. For instance, different ways of measuring employment co-vary differently with increases in the minimum wage²¹. This makes it difficult to judge overall whether the minimum wage policies have any effect.

An instance of the fifth burden in action is when trying to determine whether abortion at a given stage in the pregnancy is permissible. There are a number of competing considerations on both sides of the issue. For instance, there is a question of whether the foetus can survive outside of the womb, and if it can, what interests it may have. At the other side of the issue is what burdens the woman may face if she carries the foetus to term, versus what risks she faces in removing the foetus without injuring it, versus any risks inherent in the abortive procedure. As per the fifth burden, these considerations make it difficult to judge whether abortion at a given stage is permissible.

The contrast between the two burdens becomes clear if we note that in principle, there could be some issue which is not subject to the first burden but which is subject to the fifth. As we saw with regard to abortion, the difficulty in assessing whether it was permissible was a result of conflicting pro-tanto duties and permissions. On one side is a pro-tanto duty regarding the welfare of the foetus. On the other side is a pro-tanto permission to use defensive force against other beings which might threaten it. In principle, it is possible that the evidence we have for each

²¹ For instance, David Card and Alan Krueger's (1994) study using survey data from employers indicates that minimum wage increases have no statistically significant effect on unemployment. By contrast, David Neumark and William Wascher's (1995) analysis of unemployment over the same time period and region using administrative payroll records showed unemployment was exacerbated to a statistically significant degree.

pro-tanto duty or permission is clear and not conflicting. For instance, the foetus may be at a late stage in the pregnancy and the statistical evidence clearly indicates that it is likely to survive with medical assistance outside the womb if extracted pre-term. At the same time, statistical and diagnostic evidence might clearly show that there are significant risks to the expectant woman if she continues to carry the foetus to term or an attempt was made to extract it without killing it. Therefore, we can suppose that there are at least some cases in which the evidence is not conflicting with regards to each pro-tanto duty or permission. Therefore, in this case, the first burden is not playing a significant role in generating reasonable disagreement. However, even where the evidence is not conflicting, the existence of these conflicting pro-tanto duties and permissions still makes deciding whether abortion in the given case is permissible difficult.

Having distinguished the fifth burden from the first burden, we are now in a position to see how the fifth burden also presupposes permissivism. The claim made by the fifth burden is that with many different moral considerations on both sides of an issue, it is difficult to make an assessment and that this makes rational disagreement possible. Notably, no mention is made of whether the disagreeing parties disagree about the strength of the various considerations. That is to say, according to the fifth burden, people would rationally disagree about a given moral issue even if they agreed about what the relevant moral considerations were and the relative strengths of these considerations. This would be the case as long as there were many considerations on both sides of an issue, making it difficult to make an overall assessment. If people agreed about the relevant moral considerations and their relative strengths, then their evidence has to be relevantly similar and hence can be taken to be shared.

The point here, of course, is not that whenever there are multiple considerations on both sides of an issue, people share the same evidence. Rather, when people do not share the same evidence, we do not need to appeal to the fact that there are multiple considerations on both sides of a given issue in order to explain rational disagreement. Consider the following case:

Kitten Torture: Sid was raised in an environment where he was not exposed to suffering in either in himself or in others²². Therefore, he does not know what suffering or pain is and does not know what

²² Buddhist mythology describes Prince Siddhartha, the person who would eventually become the Buddha as having been raised shielded from any knowledge of pain and suffering.

torture is. Therefore, he does not know that torturing kittens for fun is wrong. Bud is acquainted with pain and therefore knows that torturing kittens for fun is wrong.

Kitten Torture provides us with a case in which there are not many considerations on both sides of an issue. We can therefore take it that the fifth burden is not operative in this case. If someone knows what pain and suffering are, they should know that the intentional infliction of pain and suffering on a kitten for fun is wrong. This is therefore not a situation in which multiple considerations on both sides of an issue make an overall assessment difficult. Sid lacks evidence about pain and therefore does not know what it is, while Bud possesses said evidence. This difference in evidence is sufficient to explain rational disagreement between Sid and Bud about a moral question. Given that a difference in evidence is sufficient to explain rational moral disagreement when the fifth burden is not operative, the fifth burden must presuppose that disagreeing parties would disagree even if they did share the same evidence. That is to say, the fifth burden presupposes permissivism. We can now turn our attention to the sixth or last of the listed burdens of judgment.

2.7 The Sixth Burden

According to the sixth burden:

“f. Finally, as we note in referring to Berlin’s view (V:6.2), any system of social institutions is limited in the values it can admit so that some selection must be made from the full range of moral and political values that might be realised. This is because any system of institutions has, as it were, a limited social space. In being forced to select among cherished values, or when we hold to several and must restrict each in view of the requirements of the others, we face great difficulties in setting priorities and making adjustments. Many hard decisions may even seem to have no clear answer.” (ibid)

The sixth burden firstly concerns cases in which trade-offs between certain moral and political values must be made because no social institutional arrangement can fully realise (Berlin 1988) all of them. The key point to note here is that disagreement here is purely practical. It arises even when both parties share all the same normative beliefs. It is perhaps more accurate to call this, not a source of rational disagreement as such, but one of the circumstances of politics that gives

rise to practical conflict. As such, since our concern in this thesis is about epistemically rational disagreement, the sixth burden is of little concern.

Summing up, I have clarified in this chapter what it means to share a total body of evidence and demonstrated that four of the six listed burdens of judgment presuppose permissivism. One burden, the sixth according to which “disagreement” arises because people have to make trade-offs between equally important goals, is completely irrelevant to rational disagreement as such and only one burden, the fourth, accounts for rational disagreement by appealing to different evidence. Even for this burden, of the three possible ways in which people’s differing backgrounds might generate rational disagreement, two of those ways are compatible with them having the same total body of evidence and hence presuppose permissivism. The burdens of judgment are specifications of when rational disagreement is possible. Insofar as they mostly presuppose permissivism, this gives us some reason to think that the standard version of political liberalism requires permissivism to be true. Moreover, recall that other political liberals like Larmore (2008), Gaus (2011) and Quong (2010) also regard the burdens of judgment or something very much like them as explaining the existence of rational disagreement about a great many moral and political matters. If the burdens, or anything like them presupposes permissivism, and permissivism is false, political liberalism, or at least the standard version of it is in trouble. My task in the rest of the thesis is to show that permissivism is in fact false.

Chapter 3: The Best Version of Permissivism

Background

In the second chapter, I showed that most of the burdens of judgment presuppose permissivism. In the remaining chapters, I shall hence show that permissivism is false by arguing for the uniqueness thesis according to which:

Uniqueness: Given a total body of evidence and proposition, at most one doxastic attitude towards that proposition is the rationally justified one.

Before I proceed to arguing for uniqueness, I shall clarify again the terms of doxastic attitude and evidence. In general, a doxastic attitude may refer, on a coarser grained account, to something like a belief, disbelief or suspension of judgment. On a maximally fine grained account, we may wish to represent doxastic attitudes in terms of credences or precise probability values. While not taking a stand on which is the right way to conceive of our doxastic attitudes, the arguments I shall be offering will be compatible with both coarser and finer grained analyses of doxastic attitudes.

By the term evidence, I mean anything that could bear, one way or another, on the proposition at hand. This, as I have mentioned before is compatible with different views about what counts as evidence. Evidence may, for instance, be only propositional or ultimately non-propositional, or some mixture of the two. For ease of presentation, most of the cases I will be presenting will involve propositional evidence, but that is not to be taken to mean that I endorse any particular conception of evidence as being necessarily either propositional or non-propositional. Evidence, as I have discussed previously, is evidence that is possessed by the agent and is therefore a mental state. In addition, as I have mentioned before, a mental state like a belief counts as evidence only if it carries some positive epistemic status like being justified, known or reliably formed. Nothing I will discuss in this thesis hangs on which positive epistemic status makes any given mental state count as evidence.

Finally, I should emphasise that uniqueness, as I have specified, implies that there is at most one rational doxastic attitude which is permissible, given the evidence. I do not rule out the possibility that there might be cases where there is no fully rational response to the evidence. I also do not commit myself to the claim that

evidential reasons are the only reasons for belief. I do however commit myself to the claim that if there are non-evidential reasons, and if these pull in a different direction from evidential reasons, there may be no fully rational attitude to take given the evidence.

Uniqueness can be contrasted with permissivism according to which:

Permissivism: Given a total body of evidence and proposition, it is possible that more than one doxastic attitude towards that proposition could be rationally justified.²³

Correspondingly, we might call a case in which the body of evidence justifies more than one doxastic attitude towards a given proposition a permissive case. The uniqueness thesis can then be restated as the claim that there are no permissive cases. The general strategy for arguing for uniqueness is to show that supposing that a given case is permissive requires us to accept certain implausible claims. I will do this by drawing upon a distinction made by Kopec and Titelbaum between intrapersonal and inter-personal uniqueness (Kopec and Titelbaum, 2016). Intrapersonal uniqueness restricts the permitted range of attitudes for each person. That is to say according to intrapersonal uniqueness, if some doxastic attitude A1 is rational for Bob, then no other attitude is rational for Bob. But it may be that a different attitude A2 is rational for Adam. Inter-personal uniqueness is logically stronger in that if attitude A1 is rational for Bob, no other attitude is rational for Adam or, for that matter, anyone else, given a body of evidence. In order to show that reasonable disagreement about fundamental moral principles is not possible, I need to show that inter-personal uniqueness is the case.

In this chapter, I shall argue that any plausible account of permissivism must be consistent with intrapersonal uniqueness. To start off with, I shall first present Roger White's Evidence Pointing Problem (White 2014) and his Arbitrariness Objection (ibid) to permissivism. I shall then present Miriam Schoenfield's response to White, in which she argues that the latter's arguments only show that intrapersonal uniqueness is true. Schoenfield argues that rational disagreement is possible if people have different epistemic standards (Schoenfield 2014). I shall briefly discuss an objection to the Arbitrariness Objection and an alleged counterexample to both of

²³As many authors on this topic have noted, many epistemological theories like van Fraassen's Voluntarist Epistemology, various types of Epistemic Conservatism, Foley/James style instrumentalism, Rawlsian reflective equilibrium, orthodox Bayesianism and coherentism require permissivism to be true (White 2005; Kelly 2010; Ballantyne and Coffman 2011; Schoenfield 2014).

White's arguments and argue that we can disregard them. This should show that the only plausible route to permissivism is Schoenfield's by which people can permissibly respond differently to evidence if and only if they permissibly accept different epistemic standards. This is consistent with the claim made by the second burden of judgment, according to which, people can rationally disagree because they disagree about the strength of various evidential considerations. In Chapter 2, I interpreted this disagreement about the strength of evidential considerations to refer to having different epistemic standards. The argument in this chapter will show that having different epistemic standards is not just sufficient for rational disagreement given the same evidence, it is necessary as well. Given this conclusion, in subsequent chapters, I shall present different accounts of permissivism which are consistent with intrapersonal uniqueness and argue that each account nevertheless faces insurmountable problems.

3.1 The Evidence Pointing Problem

Two classical arguments for uniqueness were made by Roger White which take the form of two objections to permissivism. The first is the Evidence Pointing Problem while the second is the Arbitrariness Objection. The basic idea behind the Evidence Pointing Problem is that evidential support is unidirectional.

"It cannot be that E supports P but also that it supports not-P. Whatever is evidence for P is evidence against not-P. If it could be that the evidence supports both P and not-P then apparently one could rationally hold both contradictory opinions at once. But that can't be right." (White 2014: 314)

Schoenfield provides an analogy with a dial. Suppose there is a dial that indicates where the total body of evidence points to, then the dial can only point in one direction. If permissivism entailed that, in at least some cases, the dial pointed in more than one direction simultaneously, then permissivism cannot be true (Schoenfield 2014: 199-200). It might seem as if a body of evidence could point in more than one direction if part of the evidence supported P and part of the evidence supported not-P. However, it would be mistaken to think that the evidence supported both P and not-P just because parts of the evidence supported one or the other. Rather, if the parts of the evidence that supported P were, together, stronger than those parts which supported not-P, then the evidence as a whole supports P. If they were weaker, then the evidence as a whole would support not-P instead. And if

the conflicting components of the evidence were of equal strength, the evidence as a whole would neither support P nor not-P.

White also argues that evidential support relations hold necessarily (White 2014: 313). By this he means that if a body of evidence E supports a proposition P, then E necessarily supports P. He provides two arguments to support this claim. The first argument is that if evidential support relations are contingent, it would be unclear how one would be able to assess what the evidence supported (ibid: 314). The thought is that, suppose there is some body of evidence E, and there is a question of whether it supports a belief B. The only way to determine whether E supports B is, either if the support relation is necessary²⁴ or if there is some further evidence E', which would indicate whether E supports B. If the question of whether E supports B is neither *apriori* nor supported by other evidence, it is unclear how we can determine whether to believe B on the basis of E. If the question of whether E supports B depends on further evidence E', we need to determine whether E' supports the proposition that E supports B. This support relation, in turn must also either be necessary or depend on yet further evidence E''. Eventually, either we run out of evidence, in which case we cannot determine whether to believe B on the basis of E, or we rely on some evidential support relation which is *apriori*. Hence, when assessing what one's total evidence supports, the only way in which such an assessment is possible is if the support relation were necessary.

The second argument White offers against the contingency of evidential support is the following: If the question of whether E supports B is determined by some contingent fact, then there are things that we may do which do not change the evidence, but which nevertheless change the fact about whether E supports B. In an example raised by White on climate change, he supposes that our evidence actually supports climate change. He argues that if this support is contingent, there may be things people could do to change whether the evidence supports climate change without changing the evidence or the climate. However, White argues that this seems absurd (ibid). Since people cannot change whether it is rational to believe a proposition based on the evidence without changing the evidence, evidential support cannot be contingent.

White's claim that the evidential support relation is a necessary one is crucial to his argument for uniqueness. If the evidential support relation holds necessarily, and a

²⁴ Here, what White means by necessary is that true facts about whether E supports B are like mathematical and logical truths which need no further evidence and are entailed by any and even the null set of propositions.

body of evidence cannot support both a proposition and its negation at the same time, then if the evidence, E, supports a belief B for one person, it supports B for everyone who possesses that evidence. Since E supports B for everyone who has E, B, and only B is justified for everyone given that their total body of evidence is E. This last claim just is uniqueness.

To be clear, the claim that evidential support is not contingent is controversial. Consider White's two arguments for evidential support being necessary. White considers it unacceptable that people may not be able to assess whether their evidence supported a proposition or that people could change whether their evidence supported a proposition without changing the evidence. However, those who reject White's necessity claim may be willing to accept these supposedly unacceptable implications. For instance, with respect to the first objection, externalists like Goldman are willing to accept that we might not be justified in believing that E supports B even when it really does (Goldman 2010: 202). On Goldman's account, he only provides the truth conditions for when a belief is justified, not an instruction manual for how to acquire justified beliefs (Goldman 1979: 90-91). Therefore, claiming that people could not figure out what they ought to believe on the basis of their evidence, if evidential support was contingent, is not an objection to the latter claim.

With respect to the second objection, a reliabilist like Goldman could also accept that there are some circumstances in which we could change whether the evidence supports our belief without changing the evidence. Consider the following case:

Seeing Red: Normally, when an object appears red to John, it actually is red. In world W1 John is currently looking at a red object. His evidence consists of his current experience of seeing an object that appears red to him as well as the cumulative of his previous confirmation of red-seeming objects to be in fact red. In world W2 everything is the same, except that just as John looks at the red object his prankster friend Mary without John's knowledge shines red light on a number of white objects just outside of John's field of vision to make them appear red.

In both W1 and W2, John has the same evidence. However, a reliabilist could plausibly claim that different beliefs are justifiable to John in W1 and W2. In W1, a reliabilist would say that John's evidence does reliably indicate that the object that he is looking at is red. In W2, the reliabilist could deny this since inferring that an

object is red on the basis that it appears red, and that previous red-seeming objects have turned out to be red does not reliably indicate that a given red-seeming object is in fact red. If so, then the fact that evidence supports something different in W2 is the product of Mary's actions. Therefore, it seems that reliabilists and other externalists about justification have reason to think that evidential support is not a necessary connection.

The fact that the non-contingency of evidential support is a controversial claim is crucial, as Schoenfield's defence of permissivism relies on the claim that evidential support is contingent. I shall discuss her account of permissivism later in this chapter. Before I do that, let me first present White's Arbitrariness Objection to permissivism.

3.2 The Arbitrariness Objection

White argues that permissivism is arbitrary by presenting the following case which I shall summarise:

Court Case: In a given court case, which I recognise as being permissive, the evidence supports both the guilty and not guilty verdicts. Perhaps I was told this by the epistemology oracle! Before coming to a conclusion by working through the evidence, I know that whichever conclusion I reach, I will have a justified belief. In that case, I could do equally well by swallowing one of two pills. One pill (GUILTY) will make me believe that the defendant is guilty while the NOT-GUILTY pill will likewise make me believe him not guilty when taken. In advance of reasoning about this permissive case, I know that reasoning correctly on the basis of my evidence is just as likely to yield me one verdict just as another. As far as the outcome is concerned, I would do no better than if I had randomly popped a belief inducing pill. In addition, I would save lots of time and effort. It would seem therefore, that as a permissivist, I should prefer popping a pill to reasoning based on the evidence. (White 2014: 315-316).

However, it is arbitrary and hence irrational to form beliefs by popping a randomly chosen pill. It is always better, qua epistemic rationality, to reason on the basis of my evidence. Therefore, whenever I suppose that a given case is permissive, I am committed to seeing my reasoning on the basis of the evidence as irrational. Ballantyne and Coffman (2011: 9) reconstruct White's argument more explicitly:

1. "If you come to believe that B is based on permissive evidence, then you think it's possible that you rationally base the opposite belief on E. [By definition of 'based on permissive evidence']
2. If you think it's possible that you rationally base the opposite belief on E, then you should think B was formed in a way no more likely to yield a true belief than is arbitrarily choosing a belief (regarding B's content).
3. If you come to believe that B is based on permissive evidence, then you should think B was formed in a way no more likely to yield a true belief than is arbitrarily choosing a belief (regarding B's content). [from (1)–(2)]
4. If you should think B was formed in a way no more likely to yield a true belief than is arbitrarily choosing a belief (regarding B's content), then B is irrational.
5. If you come to believe that B is based on permissive evidence, then B is irrational. [from (3)–(4)]" (ibid)

As mentioned earlier, insofar as an agent sees that she is in a permissive case, she is committed to regarding her reasoning on the basis of the evidence as arbitrary and hence irrational. The reason for this is that to view your reasoning as arbitrary is to view it as being driven by something other than the reasons available to you. This reconstruction of the argument by Ballantyne and Coffman can help us identify which premise any given reply to the Arbitrariness Objection rejects. Rejecting the conclusion, (5), most immediately requires rejecting either (3) or (4). Denying (4) involves asserting that, in at least some cases, it could be rational to arbitrarily choose a belief. I shall later, briefly discuss some counterexamples which purport to show that this is the case. As I shall make clear, these cases even if genuinely counterexamples to uniqueness, are not generalizable beyond very specific and narrow conditions. The more fruitful avenue to rejecting uniqueness is to reject (3). Rejecting (3) requires us to reject at least one of (1) or (2). As we can see from White's Court Case, if I think that I could be equally rational in inferring the opposite of my current belief from the total evidence I have, deliberating on my evidence is no more likely to get me the correct answer than some arbitrary process. That leaves rejecting (1). In my discussion of Schoenfield's account of permissivism, I shall show how she develops an account of permissivism, according to which it is not rational for me to believe the opposite of what I currently rationally believe, based on the same evidence. Having discussed White's two main arguments for

uniqueness, I shall present Schoenfield's account of permissivism and show how it addresses White's arguments.

3.3 Epistemic Standards

Schoenfield agrees with White that a body of evidence cannot support both P and not- P at the same time. However, her reasoning here is that the evidence does not support any proposition simpliciter; rather it can only support some proposition or other relative to some epistemic standard, provided that the standard is consistent with the principles of rationality. By epistemic standard, she means the following:

“[We] can just think of a set of standards as a function from bodies of evidence to doxastic states which the agent takes to be truth-conducive. Roughly, this means that the agent has high confidence that forming opinions using her standards will result in her having high confidence in truths and low confidence in falsehoods.”
(Schoenfield 2014: 199)

We might think of epistemic standards as consisting of sets of epistemic norms. These norms may be something like “Believe that P if it seems to you that P and you have no defeaters for P ” or they may say that “If you have experimental evidence with N samples showing result R , form attitude A towards P ”. Schoenfield's claim is that relative to a given standard, the evidence supports P , not- P or neither P nor not- P . It cannot support both. Schoenfield thus agrees with White that evidential support is unidirectional. She just denies that it is non-contingent. Rather, in her view, the evidential support relation holds relative to the agent's epistemic standards, provided that the standards do not violate what she calls the principles of rationality.

For Schoenfield, the principles of rationality are substantive rational requirements on epistemic standards. Schoenfield (*ibid*) says little about the principles of rationality, but we may nevertheless fill in some of the gaps. We might think of these principles as general constraints on doxastic attitudes that said attitudes must conform to, if they are to count as being rational. For instance, if we think that doxastic attitudes are best represented in terms of credences, we might think that the principles of rationality would include a chance-credence principle similar to that of Lewis's Principal Principle (1986). According to the Principal Principle, an agent's credence in a proposition P must match her belief in the objective chance that P is true. In addition, we might also think that certain consistency and completeness constraints

apply. Hence an agent's credences in P and not-P respectively must add to 1. Likewise, an agent's credence in the disjunction of two propositions, P or Q, must equal the sum of her credences in P and Q respectively minus her credence in their conjunction, P and Q. Or, for more coarse grained attitudes, the principles of rationality might say that an agent must not believe both P and not-P at the same time. Perhaps more strongly, the principles might require that if and only if an agent believes that P she must disbelieve not-P. With regards to induction, it might say that given a sufficiently large sample size it is permissible to generalise and believe that all Gs are Fs.

The exact content of these principles need not concern us for now, except to note that these principles constrain which epistemic standards agents may permissibly accept²⁵. Any epistemic standard which recommends doxastic attitudes that violate these constraints is not one that could be permissibly adopted. However, while these principles are substantive constraints, they are, according to Schoenfield, nevertheless very general and therefore not so strong as to fix one standard as uniquely appropriate for everyone (ibid: 202). For instance, the principles of rationality might say that testimonial evidence is important without specifying exactly how important it is. Accordingly epistemic standards which instruct agents to ignore all testimony or contrariwise to completely trust all testimony regardless of how implausible the claim, could be ruled out as violating these principles of rationality.

However, if Schoenfield is correct, the principles of rationality are not so specific as to specify exactly how much we should weight testimonial evidence in a given situation. Different epistemic standards would prescribe different ways of weighting testimonial evidence in an agent's deliberations. There may be many standards which abide by the broad constraints delimited by the principles of rationality and any doxastic attitude formed on the basis of such a standard is, on Schoenfield's picture, rational.

To further illustrate, suppose an agent, Sally's, epistemic standard includes an anti-inductive norm, according to which, she should be less confident that P will obtain in situation S, the more often she observes that P in similar situations. By complying with this norm, she becomes less confident that the sun will rise in the east the next day the more she observes that it has risen in the east in the past. Similarly, she becomes more confident that the next raven she sees will not be black the more black ravens she sees. She also becomes more confident that anti-induction will

²⁵ See Schoenfield's footnote 16 on p216.

deliver a true verdict the more often it fails to deliver a true verdict in the past. Presumably, the principles of rationality include a principle which says something like “induction is a reliable way to form beliefs”. If so, then, on any plausible account of the principles of rationality, said principles would rule out any epistemic standard which included an anti-inductive norm. The mere fact that Sally accepts such a standard is insufficient to render the beliefs she forms on the basis of that standard rational. The proviso that epistemic standards not violate the principles of rationality places limits on what would be an acceptable standard for any given agent. Within those limits, facts about what the evidence supports are relative to the agent’s epistemic standards. It is therefore possible, on Schoenfield’s account for there to be multiple permissible epistemic standards.

This makes clear how Schoenfield avoids the Evidence Pointing Problem. To draw on the analogy of the dial, Schoenfield is not claiming that the dial can simultaneously indicate both P and not-P. Instead there would be different dials; some would only point to P while others would point to only not-P. Each dial would represent an epistemic standard. Given each person’s epistemic standards, only one doxastic attitude towards a given proposition is justified on the evidence. However, given that people can permissibly hold different epistemic standards, they can still rationally disagree about a proposition, even if they have the same evidence.

As we can see, Schoenfield’s account of permissivism is consistent with intrapersonal uniqueness. We can further see the appeal of Schoenfield’s account by noting how it addresses White’s Arbitrariness objection.

The crux of White’s Arbitrariness Objection is that it would be arbitrary to form beliefs by randomly taking a pill. Schoenfield does not deny this. Her disagreement with White is about whether permissivists ought to think that forming a belief by reasoning on the basis of the evidence is better than randomly taking a pill. White believes that permissivists ought to think that reasoning on the basis of the evidence is no better than taking a pill, because by the agent’s own lights, she is just as likely to reach the truth as to believe a falsehood if she reasons from her evidence. However, on Schoenfield’s account, according to the agent’s own epistemic standards, reasoning on the basis of the evidence will be more likely to reach the truth, whereas randomly taking a pill is as likely to yield a true belief as a false one (ibid: 201).

The Evidence Pointing Problem and the Arbitrariness Objection are two powerful arguments for uniqueness. Schoenfield's account of permissivism provides a framework through which various epistemologies could address these arguments. Schoenfield deliberately says little about epistemic standards. This is so as to accommodate different accounts by which people's epistemic standards could differ. She provides some examples:

“Some people think of them as rules of the form “Given E, believe p!” Others think of them as beliefs about the correct way to form other beliefs. If you are a Bayesian, you can think of an agent's standards as her prior and conditional probability functions. (ibid: 199)”

Thus, for instance, when confronted with the Evidence Pointing Problem, the subjective Bayesian could say that given the priors she has, the evidence points to one conclusion, but that does not mean that people cannot have different priors. Similarly, when confronting the Arbitrariness Objection, the Bayesian could say that given her priors, there is only one way in which she can update her credences, given the evidence. That means that taking a random belief inducing pill would not necessarily result in the doxastic attitude she ought to have given her priors. We can see, therefore, that the subjective Bayesian endorses intrapersonal uniqueness, even if she does not endorse interpersonal uniqueness. Similarly, any permissive epistemology which can offer up a Schoenfield-style response²⁶ to White's arguments must be consistent with intrapersonal uniqueness. Permissivism can be successful only if there is some plausible account of how people may permissibly accept different epistemic standards.

If I can show that no such account is plausible, then I would be able to demonstrate that if we accept intrapersonal uniqueness, we have good reason to accept interpersonal uniqueness. However, if permissivists can reject intrapersonal uniqueness, then they could rightfully reject interpersonal uniqueness, and arguing that intrapersonal uniqueness entails interpersonal uniqueness would be pointless. Therefore, before I show why interpersonal permissiveness is not consistent with intrapersonal uniqueness, I shall explore two objections to intrapersonal

²⁶ Ballantyne and Coffman (2011: 12) specify that any epistemology according to which some other factor apart from the evidence is necessary to fix what attitude towards a proposition is rational given the evidence will be permissive. Here too, intrapersonal uniqueness is not denied. While Ballantyne and Coffman do not explicitly mention epistemic standards, given the vagueness of what Schoenfield means by epistemic standards, we can think of the additional factors that Ballantyne and Coffman refer to as those features of the situation which license different epistemic standards.

uniqueness. The first objection is Meacham's (2014) objection to the Arbitrariness Objection. The second involves two purported counterexamples to intrapersonal uniqueness; I shall then argue that we can set aside cases like these.

3.4 On the Irrelevance of Procedural Considerations to Permissivism

Meacham's counterargument to the Arbitrariness Objection is that the question of whether there are procedural constraints on epistemic rationality is orthogonal to the question of whether evidential support is permissive (Ibid: 1207-1208). After all, according to Meacham, someone could consistently affirm both that there is only one attitude supported by the evidence and that how people reach that attitude does not matter. Similarly, one could affirm that more than one attitude fits the evidence but also affirm that rationality requires that these attitudes be arrived at by a non-arbitrary process (ibid). Meacham therefore denies that if an agent sees that her evidence is permissive, she is committed to seeing reasoning on the evidence as no better than forming beliefs by some arbitrary process.

It is unclear, however, if Meacham's argument works. Suppose an agent, Sally flips a coin to decide how to respond to her evidence. By luck, she happens to pick an attitude which uniquely fits her evidence. It would not be inconsistent for an impermissivist, someone who rejects permissivism, to retrospectively regard Sally's doxastic attitude as rational just because it happens to fit her evidence. However, an impermissivist necessarily has to regard the prospect of deciding what to believe by flipping a coin as rationally suspect. The impermissivist believes that at most one of the attitudes picked by a coin would be rational. On the other hand, reasoning sufficiently well on the basis of the evidence will certainly result in a rational attitude. No arbitrary process could deliver such certainty to the impermissivist.

A corollary to this is that any process which would yield attitudes that uniquely best fit the evidence would not be arbitrary. Suppose there was a pill that, if taken, would produce in the agent doxastic attitudes on a given topic which would fit her total evidence. Whatever else is wrong with taking the pill, it certainly is not arbitrary if we know in advance that the pill will produce attitudes which fit the evidence. Similarly, if an epistemology oracle were to reliably inform us which attitudes fit our evidence, then we would have strong reason to listen to the oracle. Therefore, while an arbitrary process might yield an attitude that fits the evidence, it could only do so as a matter of chance. Agents who accept uniqueness necessarily have at least instrumental reason to avoid arbitrary ways of arriving at doxastic attitudes.

Let us also consider the other half of Meacham's argument. It is true that a permissivist could consistently affirm that a doxastic attitude is rational only if it is produced by a process of good reasoning, and believe that good reasoning alone does not fix which attitude an agent will arrive at when she deliberates on her evidence. However, the permissivist's commitment to non-arbitrary processes seems under-motivated. This is not to deny that a procedural constraint on doxastic rationality is plausible. The point is that the permissivists seem to undermine the grounds for caring about procedural constraints.

To see why, consider what I call Turri's Court Case.

Turri's Court Case: Mansour is on trial for murder. The evidence against Mansour consists of four propositions which the jury know:

P1: Mansour had a motive to kill the victim.

P2: Mansour had previously threatened to kill the victim.

P3: Multiple eyewitnesses place Mansour at the crime scene.

P4: Mansour's fingerprints were all over the murder weapon.

A jury member, Improper, reasons in the following way: The tea leaves say that P1-P4 support a guilty verdict. P1-P4 are true.

Therefore Mansour is guilty of murder (Turri 2010: 315-316).

Intuitively, we are reluctant to call Improper's belief rational even though her belief fits the evidence. This suggests that doxastic rationality requires more than just that one's attitude fit the evidence, it also requires that it be arrived at in the right way. Nothing prevents a permissivist from having this intuition in response to this case. It is, however, unclear what could ground such a commitment for the permissivist.

Consider that what seems to go wrong for Improper is that tea leaves are not an appropriate means of figuring out which attitudes fit a body of evidence. Improper seems to be lucky that she believes that her tea leaves tell her what the evidence actually happens to support. Suppose that an agent, Sally, were to encounter White's court case where the belief that the defendant was guilty was just as rational on the evidence as the belief that he was not. Once Sally reasoned her way to the belief that the defendant was guilty, she would have to consider herself lucky that she arrived at the right belief (from her perspective) since she might just as easily have reasoned that the defendant was not guilty instead. It seems that Sally faces a dilemma: Either believing the defendant not guilty is just as rational a response to the evidence as believing that he is guilty in which case Sally must view

herself as being incredibly lucky for getting it right or she is not lucky and reasoning on the basis of the evidence would not have just as easily led her to the opposite belief. However on the second option the opposite belief is not as well supported by her body of evidence as her own. In that case she cannot view it as a permissive case.

In this way, Sally's belief formation method seems objectionable for the same reasons that Improper's is. If Sally's being lucky is not objectionable for permissivists, it is unclear how they could also similarly object to Improper's reasoning. It follows, therefore, that permissivists seem to lack any reason to care about following procedural constraints when arriving at a belief.

The permissivist might still reply that while in a given case, the permissivist might indeed be lucky if she gets the right answer by reasoning on the basis of the evidence in a permissive case, but not every case is a permissive case. If most cases were impermissive, then reasoning on the basis of the evidence can be expected to get you closer to the truth on average.

However this reply is inadequate. Even if the permissivist has a reason to reason on the basis of the evidence in impermissive cases, she has no reason to do so in permissive cases. We might draw an analogy with act utilitarianism here. For the act utilitarian, the right action is the one that actually maximises the aggregate happiness. However, because it is difficult to calculate the aggregate happiness generated by each option every time the agent faces a choice, she may adopt rules of thumb that if followed can be expected to maximise aggregate happiness because attempting to calculate aggregate happiness for each decision would be counterproductive. Most of the time these rules of thumb prescribe the action which maximises aggregate happiness. However, the rule of thumb can sometimes prescribe the wrong act. The rule of thumb does not change what she ought to do, it just prescribes an effective decision procedure. Where the agent knows that the act prescribed by the rule of thumb does not maximise happiness, she ought to maximise happiness instead of following the rule. Similarly, while reasoning on the basis of evidence is better than flipping a coin on average, in permissive cases it is no better and may be worse than flipping a coin. If an agent knows that she is in a permissive case she has no reason to prefer reasoning on the basis of her evidence to flipping a coin in that particular case.

The Arbitrariness Objection seems to go through as long as there are at least some cases whereby an agent can come to know that she is in a permissive case. In

those cases the agent is committed to believing that her reasoning on the basis of permissive evidence is irrational. However, permissive evidence is evidence which justifies more than one doxastic attitude towards a proposition. If reasoning on the basis of known permissive evidence is irrational, the evidence cannot justify more than one doxastic attitude towards a proposition. This means that, at the very least, there cannot be evidence that is known to be permissive. A consequence of this is that either there are no permissive cases, or all permissive cases are covertly permissive. If the first option is true, then permissivism is false. On the other hand, if the second option is true, people are rationally committed to being systematically mistaken about the nature of the practice they are engaged in. However, this seems implausible. As I will argue in the next chapter, this is implausible because it commits us to endorsing an error theory about rationality.

Summing up, Meacham argued that the Arbitrariness Objection was mistaken because it seemed to make a procedural criticism which was orthogonal to the question of whether permissivism or uniqueness was true. I argued that the two issues were not completely orthogonal to each other. Impermissivists, I argued, had a ready reason to care about an intuitively plausible constraint on rationality. Permissivists, on the other hand did not seem to have any such reason as they were subject to the same luck as someone who did not reason well but nevertheless lucked upon an attitude that fit her evidence. Therefore, Meacham's reply to the Arbitrariness Objection is not successful.

3.5 Your Belief Makes It Real

A second objection to intrapersonal uniqueness is a class of purported counterexamples to intrapersonal uniqueness. These counterexamples involve cases where the agent's belief that P brings it about that P. Koppec (2014) and Drake (2017) each provide one such counterexample. In this section, I will discuss these counterexamples and show why they can be disregarded. Consider first Drake's counterexample which I shall label 'Positive Thinking':

Positive Thinking: Ian has a life threatening illness. "Suppose that Ian does not know whether P: he will survive the illness. He has only the following information about his prospects of survival: the illness could lead to his death, and if he does survive, the battle with the illness will render the quality of his remaining life extremely low. Knowing how grim his life might be after survival, Ian is genuinely torn as to whether he wants to survive at all. Indeed, Ian's psychology is such

that he has no reason to think that he is more likely to end up believing that P rather than not P. But Ian does know that by believing that he will survive, he will in fact secure his survival; and he knows that if he believes that he will not survive, then he will not survive.” (Drake 2017: 4901)

Drake argues that since Ian knows that whatever he believes will turn out to be true, he cannot permissibly suspend judgment. The only options remaining for Ian are to believe that he will survive or to believe that he will not. However, there is no reason to think that one of these options is more rational than the other. Moreover, there is no reason to think that either option is rationally impermissible. Drake thus concludes that both P and not-P must be rational for Ian (ibid: 4902-4903). Kopec’s case, which I shall label ‘Mind Reader’, is somewhat more explicitly science-fictional.

Mind Reader: A patient is connected to a mind-reading device and knows that she is. There is a box on the table and next to the box are some cats. A neuroscientist tells her that she (the scientist) will place a cat in the box if the patient believes that the scientist will do so and leave it empty if she believes that the scientist will leave it empty instead (Kopec 2014: 405-406).

Kopec’s argument here is similar to Drake’s. There is no reason not to believe that the cat will be put in the box as the belief is guaranteed to be true. The same can be said for the belief that the cat will not be put in the box. Suspending judgment, on the other hand, seems problematic when there is a guarantee that one’s belief will come out true (Kopec 2014: 406-407).

Seeing how these cases handle the Arbitrariness Objection and the Evidence Pointing Problem is instructive. Suppose that both Ian and the patient were given the option of forming a belief by randomly taking a pill. It is hard to deny that deliberating on the evidence is no better than randomly taking a pill. Moreover, it is a given that randomly taking a pill would be arbitrary. What the counterexample denies is the fourth premise in the Arbitrariness Objection:

“If you should think B was formed in a way no more likely to yield a true belief than is arbitrarily choosing a belief (regarding B’s content), then B is irrational.” (Ballantyne and Coffman 2011: 9)

In such cases, choosing arbitrarily is not irrational because the truth of the belief is guaranteed regardless of its content. Where choosing arbitrarily is not irrational, deliberating on evidence cannot be irrational just because it is no better than choosing arbitrarily.

Regarding the Evidence Pointing Problem, it is unclear what the evidence is supposed to be indicating. Either the evidence points at both P and not-P in this case or the evidence does not point at anything. The first option falsifies White's claim that evidential support is unidirectional in all cases. The second option takes the issue of what the evidence points to off the table and hence makes the unidirectionality of evidential support irrelevant. Whatever we end up saying about these cases, it will address the Evidence Pointing Problem one way or another.

Having seen how these self-fulfilling cases provide counterexamples against intrapersonal uniqueness, what can be said against them? The main thing to note is that these cases work precisely because they involve self-fulfilling beliefs. Suppose that we were to grant that self-fulfilling cases are permissive, this would tell us nothing about whether reasonable disagreement is possible in cases where our beliefs have no influence on the truth of the object of our beliefs. We do not think it true of our ordinary beliefs that our belief makes them true.

Notice also, that in determining whether reasonable disagreement is possible, we would be equally well served by the following modification to the uniqueness thesis:

Uniqueness(-): Given a total body of evidence and proposition, at most one doxastic attitude towards that proposition is the rationally justified one except when the agent's doxastic attitude towards the proposition influences its truth value.

Uniqueness(-) explicitly carves out an exception for self-fulfilling cases. If it is true, rational disagreement about the claims that could play a role in justifying laws and policies is still not possible since those issues do not involve self-fulfilling beliefs. We can even specify an intrapersonal version of uniqueness(-) according to which for any given agent at most one doxastic attitude is rationally justified except in self-fulfilling cases. Intrapersonal uniqueness(-) is equally well supported by the Arbitrariness Objection and the Evidence Pointing Problem. It is also the case that the exception is such a bizarre condition that our general judgments about epistemic matters take such exceptions for granted. We take for granted that when deciding what attitude to take towards a proposition, the proposition is true or false

regardless of which attitude we take towards it. Given this assumption, it makes sense to think of epistemic rationality as a matter of truth-tracking rather than truth-making. Even though the assumption that the truth of a proposition does not depend on our believing it does not necessarily hold in all cases; it holds in the vast majority of cases. The odd cases where it does not hold are so rare that it seems mistaken to think of those cases as being governed by norms of epistemic rationality at all. If this is right, then we can drop the exception in uniqueness(-) and just assert uniqueness in either its interpersonal or intrapersonal versions.

It seems clear then that intrapersonal uniqueness is true, at least in all the non-trivial, non-science fictional cases. It follows that any plausible account of permissivism will have to endorse intrapersonal uniqueness²⁷. More specifically, relative to an epistemic standard, there is at most one rational doxastic attitude to take with respect to a proposition, given the total body of evidence.

As mentioned earlier in Chapter 2, Schoenfield's claim that people can rationally disagree, given the same total body of evidence, because they have different epistemic standards explains how the second burden of judgment is a source of rational disagreement. Recall that according to the second burden of judgment, people rationally disagree because they differ in the weight they assign to the various considerations for or against a proposition. This was interpreted to mean that people rationally disagreed because they had different epistemic standards. As we have seen so far, in this chapter, any plausible account of permissivism must have some account of how people could have different epistemic standards. This is consistent with the claim made by the second burden of judgment. In fact, it is even stronger: People could rationally disagree about a proposition if and only if they could permissibly have different epistemic standards

One of the central claims in this thesis is that the burdens of judgment, which specify the conditions that make reasonable disagreement possible, presuppose permissivism. In showing that permissivism is false, I show that the standard account of how reasonable disagreement is possible requires significant revision. Pursuant to this, I need to show that insofar as a burden of judgment presupposes permissivism, the disagreement caused by that burden is not rational. With regards to the second burden, given how the possibility of people having different epistemic standards is crucial to them rationally disagreeing, showing that the second burden

²⁷ In addition to Schoenfield, Ballantyne and Coffman (2011), Kelly (2014) and Meacham (2014) also endorse intrapersonal uniqueness.

is not a source of rational disagreement involves showing that people cannot have different epistemic standards. This raises the question as to how it is permissible for people to have different epistemic standards.

Permissivists have two options available to them. The first option is that for any given agent, there are two or more epistemic standards which are equally well suited for her. The second option is that each person has only one epistemic standard that is most suited for her, but the standard that is most suited for one person, may not be the standard most suited for another. I shall argue in the next chapter that if permissivists take the first option, they can still be vulnerable to the Arbitrariness Objection and the Evidence Pointing Problem. In the subsequent chapters, I shall consider various proposals that are consistent with the second option and argue that none of them will work either. This will show that if intrapersonal uniqueness is accepted, we are required to accept inter-personal uniqueness as well.

Chapter 4: Intrapersonal Uniqueness about Epistemic Standards

Background

In Chapter 3, I argued that at least intrapersonal uniqueness is true. The two main arguments for this were the Arbitrariness Objection and the Evidence Pointing Problem. The permissivist's reply to these arguments would be that at most one doxastic attitude is rational given the evidence relative to an epistemic standard, but people could have different epistemic standards. This reply mirrored Rawls's second burden of judgment according to which people could rationally disagree about a proposition if they disagreed about the strength of an evidential consideration. People's disagreement about the strength of a given evidential consideration can plausibly be thought of as reflecting their possession of different epistemic standards. In questioning how people could permissibly have different epistemic standards, two possible answers can be offered. The first of these answers, which I shall call intrapersonal permissivism about standards, is that for any given agent, more than one epistemic standard is rationally permissible for her. Thus on this view, an agent, Sally, who could permissibly make inferences on the basis of a standard S1 could also permissibly make inferences on the basis of S2. The alternative answer, which I shall call intrapersonal uniqueness about standards, is that for any given agent only one epistemic standard is rationally permissible for her. On this view, if Sally permissibly reasons on the basis of S1, then no other epistemic standard is appropriate for her. This, however, does not preclude that some other standard S2, but not S1, is appropriate for another agent, Susan. The subject of this chapter is the former view. I shall argue that both the Arbitrariness Objection and the Evidence Pointing Problem can be extended against intrapersonal permissivism about standards. If these arguments succeed, then the only viable accounts of permissivism are those which are consistent with intrapersonal uniqueness about epistemic standards.

4.1 Extending the Arbitrariness Objection

To recap, the crux of the Arbitrariness Objection was that if evidence was permissive, reasoning on the basis of evidence would be no better than forming beliefs by taking a random pill or flipping a coin. However, since these ways of forming beliefs are arbitrary, they are irrational. Forming beliefs on the basis of permissive evidence is therefore irrational. Schoenfield's response was to argue that permissivists would not choose to arbitrarily form beliefs because beliefs formed by such processes may not be the ones justified by the evidence given their

own epistemic standards (Schoenfield 2014: 201). This response seems to assume that the belief changing pill changes beliefs without changing a person's epistemic standards. A natural way to extend the Arbitrariness Objection is to modify the case such that the pill does not change people's beliefs without also changing people's epistemic standards (White 2005).

If the pill changed our beliefs by changing epistemic standards to one that is also just as rational, then we could not object to taking the pill on the grounds that the resulting belief violated our epistemic standards. After all, after taking the pill, our beliefs would still match our epistemic standards because we would have new standards. It seems as if permissivists should have no objection to taking the pill. However it is arbitrary to form beliefs by popping pills. It also seems arbitrary from an epistemic point of view to prefer our own standards just because they are our own or for any other non-epistemic reasons²⁸.

Schoenfield anticipates this objection and argues that the permissivist could still reject taking the pill as the new standards would, from her current perspective, be less truth-conducive than her current standards (Ibid). To be clear, since the permissivist thinks that her own standard is more truth-conducive than any other, her preference for her own standard is not arbitrary. However, it is unclear how the permissivist can think that one standard is more truth-conducive than any other and still think that all of these standards are equally permissible ones.

One option, as Simpson (2017) notes, is to argue that one's own standard is the most truth-conducive, but that others are good enough. However, as Simpson argues²⁹, if there is some way in which standards which fall short of being the best are less truth-conducive than the best standard, it is hard to see why adopting those other non-best standards is not a mistake and therefore irrational.

Schoenfield tries to clarify how one's own standard could be more truth-conducive than others but yet still be as rational as some others:

“We might cash this out by thinking that the principles of rationality are going to be general: they will be principles about what kinds of considerations count in favour of what kinds of hypotheses. But these sorts of general considerations are not sufficiently robust to pin down

²⁸ Simpson argued for this second aspect of arbitrariness in his response (2017) to Schoenfield.

²⁹ Although Simpson makes this argument with respect to doxastic attitudes, the argument seems equally applicable to standards.

a unique doxastic state given by any body of evidence. So even if Sally and her friend both conform to these principles, their standards may differ with regard to how exactly they weigh the different considerations and thus, in any given case, Sally and her friend might rationally come to different conclusions.” (Schoenfield 2014: 202)

The thought here is that the principles of rationality are supposed to be substantive, but general constraints on rationality. They may include Lewis’s Principal Principle, or Kolmogorov’s probability axioms or general principles claiming that induction is reliable. They may tell us that some considerations matter, but fail to specify exactly how much they matter. For instance, they may say that induction is reliable, but fail to specify exactly how many instances we are supposed to observe before we are permitted to make a generalisation. Since the principles do not tell us how much weight ought to be given to any given consideration, any epistemic standard which employs these considerations satisfies these principles and is thus rational. Some ways of weighing these considerations are more truth-conducive than others, but the principles are too general to tell us which one is the most truth-conducive.

There are two crucial claims that Schoenfield relies on. The first is that it is rationally permissible to pick any of the epistemic standards which comply with the principles of rationality. The second claim she makes is that Sally can simultaneously accept that there are different permissible ways of weighing the various considerations, and that her own epistemic standard is the most truth-conducive.

I shall offer an objection pertaining to each claim. Regarding the first claim, I shall argue that if the principles of rationality are insufficient to fix on one epistemic standard, epistemic standards would seem to be arbitrary. Regarding the second claim, I shall argue that Schoenfield faces a dilemma: On the first horn of the dilemma, Sally believes that her own standard is the most truth-conducive, in which case she accepts uniqueness. Alternatively, she accepts that the truth-conduciveness of different standards are either equal or incommensurable, in which case, her adoption of her own standard is vulnerable to the Arbitrariness Objection. I shall discuss each objection in turn.

The first objection is that if we think that if the principles of rationality are so general and broad that they cannot fix on any one epistemic standard as best encapsulating them, then an agent’s selection of the more specific standard seems arbitrary. To see why this is the case, suppose that it is the case that the principles of rationality

are unable to pick out any given epistemic standard, but pick out a set of standards, $\{S1, S2\}$ as being standards that comply with the principles, but there is no determinate criterion for picking one of them as better. S1 and S2 pick out attitudes A1 and A2 respectively in response to the evidence E. Suppose Sally chooses S1 over S2. Either she has epistemic grounds, i.e. reasons related to epistemic rationality, to think that S1 is better, or she has non-epistemic grounds or she has no grounds to choose S1 over S2. Let us consider the last option first.

4.1.1 Choosing For No Reason

Suppose Sally chooses S1 over S2, but does so for no particular reason. This would seem to be a clear case in which Sally's choice is arbitrary and therefore irrational. The permissivist might, at this point, reply that not all arbitrary choices are irrational, some may simply be a-rational, beyond rational assessment. After all, whichever standard is chosen, one is guaranteed to be in conformity with the principles of rationality and not choosing would violate the principles of rationality. The thought here is that by not accepting a standard, one is unable to form any attitudes whatsoever on the evidence even in simple cases and that should be a violation of the principles of rationality if anything is. If Sally has no choice but to choose one of S1 or S2, then it cannot be irrational for her to arbitrarily choose one or the other. Therefore it is permissible for her to plump for one set of epistemic standards.

The permissivist's reply to this extension of the Arbitrariness Objection depends on the principles of rationality not containing any symmetry principles. A symmetry principle is any principle that picks out some unique doxastic attitude whenever two or more doxastic attitudes towards a proposition or two or more alternative propositions are equally well supported by the evidence. Two candidate symmetry principles are salient here. The first is the principle of insufficient reason³⁰. The principle of insufficient reason is a principle of reasoning that recommends that we treat evidentially symmetrical alternatives as equally likely. For instance, it requires that we suspend judgment about whether a coin will land heads if we know nothing about the coin. The reasoning is that if my evidence supports the proposition that the coin will land heads exactly as much as it supports the proposition that it will land tails, then each proposition is equally likely. Similarly, such a principle also requires us to assign a credence of 1/6 to a six sided die showing 4 when it lands. As with the coin toss, since my evidence supports the proposition that the die will

³⁰ See Howson and Urbach (2005) and Meacham (2014) for criticisms of the principle of insufficient reason. See White (2010) for a defence of the same.

show 4 as much as it supports the proposition that it will show 3 and so on, the likelihood of any one face showing must be the same as the likelihood of some other face showing. Therefore, I ought to have a credence of $1/6$ that the die will show 4.

Instead of the principle of insufficient reason, the principles of rationality might instead have what I call the Fuzziness principle.

Fuzziness: If two or more precise doxastic attitudes, A_1, A_2 to A_n , towards a proposition are equally well supported by the evidence for an agent and no precise attitude is better supported, she ought to adopt a fuzzy or imprecise doxastic attitude ranging over A_1 to A_n .

Fuzziness instructs agents to adopt imprecise doxastic attitudes³¹ in order to match the lack of precision inherent in the evidence. For instance, if a credence of 0.7 and a credence of 0.8 are both well supported by the evidence, Fuzziness requires agents to adopt an imprecise credence of $[0.7, 0.8]$ instead of either 0.7 or 0.8. The attitude $[A_1, A_2]$ can be thought of as covering both A_1 and A_2 without specifically being either A_1 or A_2 .

If either the principle of insufficient reason or Fuzziness is correct, then the principles of rationality are always strong enough to pick out one epistemic standard as correct. I shall explain how this can be for each candidate symmetry principle.

Consider, first, the principle of insufficient reason. If an agent, Sally, applies the principle of insufficient reason, instead of adopting epistemic standards S_1 or S_2 , she should adopt an intermediate standard S^* . Recall that an epistemic standard is constituted by epistemic norms and each epistemic norm is or entails a rule of the form "if your evidence is E , form doxastic attitude A towards proposition P ". By stipulation, S^* picks out an attitude A^* which lies between A_1 and A_2 . The key assumption here is that there is a doxastic attitude A^* that lies between A_1 and A_2 . This assumption is most plausible if doxastic attitudes are represented by sharp credences. For instance, consider the following case:

Impeachment: Susan and Sally are both considering whether Donald Trump should be impeached. Both of them watch the same programmes on Fox News and MSNBC. The credences 0.75 and 0.76 are both equally well supported by Susan's and Sally's

³¹ See White (2010) for a criticism of imprecise credences and van Fraassen (1990) and Chandler (2014) for a defence of imprecise credences.

evidence. Sally adopts a credence of 0.75 while Susan adopts a credence of 0.76.

In principle, it seems that an agent could have credence 0.755 towards the proposition that Donald Trump should be impeached. If some intermediate credence is always available then there is, in principle, some epistemic norm which requires the agent to form the intermediate attitude 0.755 in response to her evidence. Since epistemic standards are constituted by epistemic norms, different combinations of epistemic norms will give rise to different epistemic standards. It follows, that if there is some intermediate doxastic attitude available to an agent, there is a corresponding epistemic standard which if adopted, would require the agent to adopt that intermediate doxastic attitude. In short, where doxastic attitudes are maximally fine grained, there is always an intermediate epistemic standard S^* between two standards S_1 and S_2 . To be clear, the mere existence of the intermediate standard S^* does not suffice to make it preferable to the standards S_1 and S_2 . Rather, or so I shall argue, S^* is preferable to S_1 and S_2 if the principle of insufficient reason is true.

The reason for this is that epistemic standards encode assumptions about how to weigh conflicting pieces of evidence against each other. Applying the principle of insufficient reason requires us to moderate our confidence in these assumptions. This in turn will yield a standard that generates an intermediate doxastic attitude. I shall illustrate this point by returning to the Impeachment case.

In Impeachment, Susan's and Sally's total body of evidence consists of two sources of information: Fox News and MSNBC. Suppose that we can score the trustworthiness of news stations from 0 to 10. Epistemic standard S_1 encodes the assumption that MSNBC has a trustworthiness score of 8 and that Fox News has a score of 2. Given these trustworthiness scores, the appropriate credence for the proposition that Donald Trump should be impeached is 0.75. By contrast, S_2 encodes the assumption that MSNBC has a trustworthiness score of 9 and Fox News has a score of 1. Given these scores, the appropriate credence for the claim that Trump should be impeached is 0.76. However, both assumptions cannot be true and given uncertainty about which assumption is true, the principle of insufficient reason recommends that Sally assign the likelihood of each assumption being true a value of 0.5. The expected trustworthiness of a given News station is the trustworthiness of the station given the first set of assumptions multiplied by the likelihood of that assumption added to the trustworthiness given the second set of

assumptions multiplied by that assumption's likelihood. The expected trustworthiness of MSNBC will be 8.5 and the expected trustworthiness of Fox News will be 1.5.

Given these scores, the appropriate credence that Trump should be impeached is 0.755. Since MSNBC advocates for Trump's impeachment and is likely to present evidence that favours impeaching him, the more trust an agent places in MSNBC, the more confident she should be that Trump should be impeached. Likewise, since Fox News advocates against Trump's impeachment and tends to present evidence that favours not impeaching him, the more trust an agent places in Fox News, the less confident she should be that Trump should be impeached. This implies that the credence that is appropriate given the expected trustworthiness of the respective news sources is in between the credences given S1 and S2 respectively. Since S* requires the agent to have credence 0.755 in the proposition that Trump should be impeached, Susan and Sally should employ epistemic standard S* instead.

Summing up, epistemic standards consist of norms which govern belief formation at least in part by determining which considerations should matter more when they conflict with each other. This means that any epistemic standard must encode assumptions that explain why the considerations which the standard takes to be stronger are in fact stronger. If the principle of insufficient reason is applied, the agent applies an equal probability to each alternative assumption. The expected strength of each consideration will therefore be of an intermediate value as compared to the comparatively more extreme values encoded in the various assumptions. Correspondingly, the norms that these expected strengths encode will recommend an attitude that is intermediate between the ones recommended by the initial set of standards. Therefore, if the principle of insufficient reason is applied, agents should pick S* instead of S1 or S2.

I have, so far, assumed that there is always some intermediate doxastic attitude available. However, on more coarse grained accounts of doxastic attitudes, it might seem less plausible that there is always some intermediate doxastic attitude. Suppose, for instance that in a given case, believing that P and suspending judgment that P were equally well supported by the evidence. If there was no intermediate attitude between believing and suspending judgment, then the principle of insufficient reason, if applicable, could not possibly require the agent to pick an intermediate epistemic standard S* because no such standard would be available. It

does not follow that the principle of insufficient reason has no application. To see why, consider the following variation on Impeachment.

Impeachment II: Susan and Sally are both considering whether Donald Trump should be impeached. Both of them watch the same programmes on Fox News and MSNBC. The attitudes suspending judgment and belief are both equally well supported by Susan's and Sally's evidence. Sally adopts standard S1 and suspends judgment while Susan adopts standard S2 and believes that Trump must be impeached.

Here, we can still score the trustworthiness of each news station between 0 and 10 and as in Impeachment, Sally assumed that the trustworthiness of MSNBC and Fox News are 8 and 2 respectively. Likewise, Susan assumed that the trustworthiness of the news stations are 9 and 1 respectively. Given the principle of insufficient reason, each assumption is just as likely as the other at 0.5. This makes the expected trustworthiness of the stations turn out to be 8.5 and 1.5 respectively.

Since there is no intermediate doxastic attitude, we cannot assume that each epistemic standard encodes very specific assumptions about the trustworthiness of each news station. Instead, we must suppose that there is some threshold such that Sally's assumption falls at or below the threshold thus making suspending judgment appropriate for her. Likewise Susan's assumption falls above the threshold thus making belief that Trump should be impeached appropriate for her. For instance, if the difference in the trustworthiness of the stations is 6 or below the agent should suspend judgment while if the difference in trustworthiness is more than 6, she should believe that Trump should be impeached. If the threshold is set at 6, then given that the expected difference in trustworthiness is 7, Sally and Susan should both adopt S2 and believe that Trump should be impeached. Wherever the threshold is placed, the expected difference in trustworthiness is going to fall on one side or the other of that threshold. Insofar as the principle of insufficient reason applies, Sally and Susan cannot permissibly believe just one of the assumptions; they must take them to be equally likely. If they do so, there is only one doxastic attitude which would be justified for them given the evidence.

Summing up, even if there is no intermediate epistemic standard, when the principle of insufficient reason is applied, only one of those standards becomes permissible for the agents. Whereas the connection between the principle of insufficient reason

and the appropriate epistemic standard is somewhat indirect, the connection between Fuzziness and imprecise attitudes is simpler.

To see why, suppose that Fuzziness is true instead. Consider Sally's situation again in Impeachment. Fuzziness requires that her doxastic attitude should be imprecise $[0.75, 0.76]$. Therefore, instead of choosing S^* an intermediate standard, she should choose an indeterminate standard S that picks out doxastic attitude $[0.75, 0.76]$ ³². Any other epistemic standard would violate Fuzziness.

We started off with the claim that the principles of rationality had some amount of slack such that more than one epistemic standard was consistent with the requirements of those principles. As I have just argued, if either the principle of insufficient reason or Fuzziness is true, then the principles of rationality do not have any slack. The permissivist's claim that there is more than one rationally permissible epistemic standard can go through only if both the principle of insufficient reason and Fuzziness are false. However, we have good reason to believe that at least one or the other is true.

One key motivation for Fuzziness is the following intuition: If the principles of rationality, taken together with the evidence, are genuinely indeterminate to some degree, it is unclear why it would be epistemically permissible to be more precise than the evidence and the principles of rationality warrant. However, adopting either S_1 or S_2 would do precisely that. S_1 and S_2 each go beyond what the principles of rationality say in one way or another. S_1 and S_2 make claims about which considerations matter more, and how much more they matter. If Sally were to choose one or the other standard and hence choose between claims about which considerations matter more for no reason at all, doing so would be arbitrary and therefore irrational.

In addition, there seem to be a number of cases in which our intuitions about which attitudes are appropriate presuppose that one or the other principle must be true. Consider, for instance, the following case:

Urns: There are 11 urns in front of Sally, each with 10 balls inside. Each urn has a different mix of white and black balls including one urn in which all the balls are black and one in which all the balls are white. No two urns have the same number of black balls in them. Sally is unaware of how many black balls are in any given urn, but

³² I'm indebted to Simpson's discussion (2017) of Schoenfield's argument for this point.

she knows that each urn has a different mix of black and white balls and that each urn has a total of 10 balls. Susan reaches into an urn and picks a ball. What should Sally's confidence in the proposition that Susan will pick a black ball be?

There are two potential answers to the question. One answer is that Sally should have an imprecise confidence of $[0, 1]$ that the ball will be black. This is because she knows that the urn that Susan draws the ball from could have anywhere between 0 and 10 black balls. Given that Sally lacks information about how Susan is picking the urn or drawing balls from the urn, she knows that the objective chances that the ball will be black is anywhere between 0 and 1 and no objective chance is more likely than any other. Therefore, given some plausible chance-credence principle, Sally knows that any given precise confidence from 0 to 1 is equally well supported by the evidence. Therefore to account for this slack, Sally should adopt a maximally imprecise doxastic attitude. Notably a maximally imprecise attitude is justified for Sally only if Fuzziness or some principle that is very much like it was true.

Alternatively, it might be the case that Sally should have a confidence of 0.5 that Susan's ball will be black. This is because across all the urns, exactly half the balls are black. Sally might reason that if Susan were to wave her arms in such a way as to randomly pick out an urn and randomly pick a ball from that urn, she would in effect be picking a ball randomly from a set of 110 balls, half of which are black. If a confidence of 0.5 is required in the latter case, it must be also appropriate in the former as well. This is because Susan randomly waving her arms about cannot provide any information about whether the ball she picks will be black. Notice that Sally's reasoning here vindicates the principle of insufficient reason. Given that each urn can contain anywhere between 0 and 10 balls, the expected number of black balls in a randomly selected urn is 5. This result can only be arrived at if Sally assigns an equal likelihood to selecting each urn and, for each urn, an equal likelihood of selecting each ball. Assigning equal likelihoods like this just is to apply the principle of insufficient reason.

Notably, the question as to which doxastic attitude is permissible for Sally depends on which principle is true. If the principle of insufficient reason is true, then Sally should have credence 0.5 that the ball Susan picks will be

black. If it is false and Fuzziness true instead, she should have an imprecise credence $[0, 1]$. Either way, it is implausible that Sally could permissibly have any other credence about what colour the ball would be. However if both the principle of insufficient reason and Fuzziness are false, then there is no reason why one of those other doxastic attitudes would not be permissible. If the above is right, then either the principle of insufficient reason or Fuzziness must be true. If either of the principles is true, then there is no slack in the principles of rationality. As such there is no sense in which the principles of rationality run short and underdetermine which epistemic standard is permissible. If so, then agents cannot permissibly plump for one or other epistemic standard.

4.1.2 Choosing for Practical Reasons

Suppose that instead of choosing S1 for no reason at all, Sally chose S1 on practical or non-epistemic grounds. One possible type of practical consideration may be the stakes at hand. The views that I am alluding to are called contextualism³³ and subject sensitive invariantism. On the contextualist account, whether an agent's belief is justified depends on the context of the assessor. For instance, the question whether Sally's belief that P is justified depends on whether the person making the assessment of Sally's rationality is Susan or some other person Mark. Subject sensitive invariantism is the view that the question of whether an agent's belief is justified depends on the context of that agent. Hence, the question of whether a belief that P is justified for Sally depends on Sally's context. These views have been advanced primarily as a reply to sceptical worries raised by the possibility of evil demons, 'brains in vats' and painted donkeys. The thought here is that in the philosophy classroom, we have to take seriously the possibility that an evil demon is deceiving us about the external world and hence are not justified in our beliefs about the external world. However, outside the philosophy classroom, our context changes and we are justified in our ordinary everyday beliefs. While most of such attempts are targeted at knowledge, some recent versions of contextualism or subject sensitivity (Hambourger 1987; Wedgwood 2008; Davis 2015) are about justified belief. A further distinction is between intellectualist versions in which the context of justification or knowledge attribution is determined by which alternatives are made salient in the conversation, and more pragmatist versions where the context of justification is determined by the practical

³³ See DeRose (1992) and Cohen (1999) for defences of contextualism about knowledge.

stakes at issue. For instance, Hambourger (1987) argues that that we are justified in asserting or believing that P, if we can prudently neglect the possibility of being wrong. Thus, generally when the consequences of being wrong are more severe, it would be prudent to neglect that possibility only if the probability of being wrong was really low. That is to say, that degree of evidential support required to justify belief is very high. Similarly if being wrong is of little consequence, then it would be prudent to neglect this possibility even if there was some significant probability of being wrong. Therefore, the amount of evidence needed to justify belief is proportionally lower.

To see why, the following two cases can be compared³⁴. The first one is called Low Stakes.

Low Stakes: I wish to deposit a cheque, but there is no particular urgency with which I need to do so. I consider going to the bank today. My largely reliable but still imperfect memory tells me that the bank closes at 8pm tonight. On the basis of my recollection, I believe that the bank closes at 8pm tonight.

Contrast this with the next case which I shall refer to as High Stakes:

High Stakes: I wish to deposit a cheque, but I need the cheque to clear in the next 3 days. Therefore I need to deposit it before the bank closes today. My largely reliable but still imperfect memory tells me that the bank closes at 8pm tonight. On the basis of my recollection, I suspend judgment that the bank closes at 8pm tonight and check the internet again to confirm the opening hours.

Plausibly, the higher stakes in High Stakes rightly makes me more cautious about coming to a belief on the basis of my imperfect memory. To be clear, I am not arguing for contextualism. Rather, my aim is to sketch out an at least half-way plausible account according to which practical considerations can affect which epistemic standards an agent may permissibly apply. The thought here is that contextualism about justification might ground some version of permissivism. Practical stakes could affect the relative extents to which an agent values acquiring truths and avoiding falsehoods or her degree of epistemic risk aversion. For instance, if the consequences of acquiring a false belief about a proposition are particularly dire, but the benefits of getting it right are relatively mild, it might be

³⁴ These cases are based on DeRose's bank cases (1992).

more appropriate to have an epistemic standard which recommends suspending judgment given a body of evidence, than one which recommends belief in that proposition. As such, there might be practical stakes which give Sally reason to choose S1 over S2. Thus, Sally's adoption of S1 here can be rationally permissible in the right circumstances.

In response to this argument, I shall make the following objection: If Sally's practical reasons, together with the principles of rationality are sufficient for S1 to be a better standard for her than S2, then, she cannot permissibly adopt S2 rather than S1. To see why this would be the case, suppose Sally's practical reasons, together with the principles of rationality gave her sufficient reason to adopt standard S1 and also gave her sufficient reason to adopt S2 instead. Then her choice of using S1 rather than S2 would be arbitrary in the same way that it would have been if she had no reason to adopt S1 at all. Her practical reasons and the principles of rationality together would still have been too weak to fix S1 as the best epistemic standard for Sally. As would have been the case if Sally had adopted S1 for no reason, adopting S1 when her practical reasons are indeterminate between S1 and S2 require her to go beyond any reasons she has available to her. In doing so, she may as well have flipped a coin or popped a pill. Hence her use of epistemic standard S1 is arbitrary and therefore irrational.

The only way this arbitrariness is avoided is if Sally's practical reasons fix S1 as being best for her. Similarly, even if Susan's practical reasons are different, the only way she can avoid the charge of arbitrariness in choosing S2 is if her practical reasons fix S2 for her. Note, however, if for each agent their practical reasons together with the principles of rationality fix which standard is appropriate for her, then more than one epistemic standard is not permissible for her. While this does not preclude all forms of permissivism, this does preclude versions of permissivism which are consistent with intrapersonal permissivism about standards. In the next chapter, I shall discuss versions of permissivism according to which for each agent, only one epistemic standard is rationally permissible.

4.1.3 Choosing for Epistemic Reasons

The above discussion about how practical reasons make Sally's adopting S1 permissible can also be extended to epistemic reasons. Firstly, we can rule out the possibility that these epistemic reasons do not fix S1 for Sally. If they did not, Sally's choosing S1 over S2 would go beyond what her reasons permitted. To avoid making Sally's choice arbitrary, these epistemic reasons must fix S1 for Sally.

Given that these reasons must fix S1 for Sally, either these reasons are available to Susan and everyone else or they are not. If they are available, then these reasons, whatever they are, can be considered part of the principles of rationality. In that case the principles of rationality do fix S1 as the best epistemic standard for everyone. If, these reasons are not available to Susan, then even though some version of permissivism is still viable, that version is not compatible with intrapersonal permissivism about standards. S2 would correspondingly be the best standard for Susan. One such version of permissivism is the one offered by Simpson according to which epistemic standards vary according to cognitive capacities (2017). I shall discuss this account later in the next chapter.

Summing up, I have argued that, if the principles of rationality, together with any other considerations which would bear on the choice of epistemic standards, permit for any given agent, more than one standard, the agent's use of that standard will necessarily be arbitrary and therefore irrational. I have argued that the only alternatives are to endorse uniqueness or endorse a version of permissivism consistent with intrapersonal uniqueness about standards.

4.1.4 Viewing Alternative Standards as also Rationally Permissible

To recap, Schoenfield attempts to explain how an agent, Sally, could claim that her own epistemic standard was the most truth-conducive, but was just as rational as other epistemic standards. Her argument consists of two claims. The first is that the principles of rationality are general and do not fix one particular epistemic standard as the best. I have just addressed this first claim (4.1.1 - 4.1.3). Schoenfield's second claim is that Sally can simultaneously accept both that S1 is the most truth-conducive standard and also that Susan can rationally believe that S2 is. I shall argue that this is not possible.

I shall be arguing that there are two possible perspectives from which Sally can evaluate Susan's beliefs and epistemic standards. However, from neither perspective can Sally both endorse permissivism and the claim that her own standard is the most truth-conducive. Sally endorsing both claims involves a level confusion between these two perspectives. To see why, we need to clarify what, on Schoenfield's account, Sally believes about Susan. Sally who uses standard S1 would have the following beliefs about Susan who uses S2:

- A. Susan's inferences on the basis of S2 are rational
- B. S2 is not maximally truth-conducive

- C. It is rational for Susan to believe that S2 is maximally truth-conducive
- D. It would not be rational for Sally to believe that S2 is maximally truth-conducive

Claim A follows from the assumption that Sally accepts permissivism. To accept permissivism just is to accept that someone who uses some other epistemic standard can be rational. B follows from Sally's belief that her own standard S1 is the most truth-conducive. If S2 was maximally truth-conducive, S1 would not be the most truth-conducive because it is at best just as truth-conducive as S2. Sally needs to believe C because denying C would introduce an untenable asymmetry between hers and Susan's situations. From Sally's own point of view, her inferences on the basis of S1 are rational either because S1 is actually the most truth-conducive or at least because she rationally believes that S1 is the most truth-conducive. If Sally were to reject C, then from her perspective, neither is S2 maximally truth-conducive nor is it rational for Susan to believe so. If S2 is not maximally truth-conducive and if it is not rational for Susan to believe that it is, it is unclear why Susan's inferences on the basis of S2 are rational. Sally would have to think that the explanation for why Susan's inferences on the basis of S2 are rational is very different from the explanation for why her own inferences on the basis of S1 are rational. However, such a belief cannot be sustained. One difficulty involves accounting for why the considerations that trump truth-conduciveness for Susan do not do so for Sally. Even if Sally can provide such an account, she cannot also account for the intuition that reasoning on the basis of S2 when she is in a position to know that S1 is more truth-conducive is a serious mistake in reasoning. Given that this asymmetry about the explanation for why Sally's and Susan's inferences on the basis of their own standards are rational is untenable, Sally has to accept C. Sally needs to assert D in order for her to avoid the Arbitrariness Objection. If it would be rational for Sally to believe that S2 was maximally truth-conducive, why would she object to popping a pill that would change her epistemic standards to S2? If she cannot object to popping a pill, her reasoning must be arbitrary.

Sally, therefore, seems to have good reasons for A-D. However, it seems inconsistent for Sally to believe A-D. To see why, consider C and D: On the surface, C and D seem consistent with each other. This would be because for Schoenfield, the truth-conduciveness of any given epistemic standard can only be evaluated relative to that standard (Schoenfield 2014: 202). That means that Sally believes C because she believes C*:

C*: It is rational relative to S2 to believe that S2 is maximally truth-conducive

And similarly, she believes D only because she believes D*:

D*: It is not rational relative to S1 to believe that S2 is maximally truth-conducive.

Given that Sally can only reason on the basis of S1, there's no reason for her to believe that C* implies C. As far as she is concerned, Susan is using the wrong standards. As Simpson (2017) notes, since she can only make epistemic evaluations from the standpoint of her own standards, the fact that other standards are acceptable from the standpoint of those standards does not matter if, by her own lights, they are not maximally truth-conducive. If Sally were to try to assert the validity of other standards on the grounds of them being acceptable so long as one accepted those standards, she would be rejecting the idea that epistemic evaluations can only be made on the basis of one's own standards.

Rejecting this allows Sally to accept C and D, but requires her to reject B. To see why, we need to examine why we couldn't tell which standards are more truth-conducive independently of our epistemic standards. Consider a set of 20 propositions, 10 of which are true. Ideally, we would wish to have the highest confidence in the true propositions and the lowest confidence in the false ones. If any epistemic standard S could deliver this result, there would be no question, regardless of which other standard we initially accepted, as to whether S was the rational standard to accept. However, not only is it unlikely that there is any such standard, our evidence is rarely so complete or strong that they conclusively indicate the truth of every question we may happen to be interested in. As a result, sometimes, we may end up having a high confidence in a falsehood and a low confidence in a truth. Between two standards S1 and S2, if S1 delivers a higher confidence in more truths and a lower confidence in more falsehoods than S2, then there is no question as to which standard is better. However, not all standards can be ordered so neatly that there are no trade-offs to be made between them. Suppose that instead of S1 being superior to S2 along all dimensions, S2 delivers lower confidences in falsehoods at the cost of lower confidences in truths as well. Or alternatively, it delivered higher confidences in truths at the expense of higher confidences in falsehoods. There seems to be no way, apart from adopting some scoring system, of measuring which standard is more truth-conducive.

However, there are many possible scoring systems. Some may of course deliver results that are obviously absurd³⁵, but there may still be many scoring systems which provide different plausible ways of evaluating the overall truth-conduciveness of a standard. Each scoring system will select one standard or another as the one which is the most truth-conducive. The question of which standard is the most truth-conducive thus depends on which scoring system we have most reason to adopt. However, the reasons we have for choosing one scoring system over the other are that they do a better job of identifying which inferences are rational given the evidence. This means that which scoring system we adopt is going to depend on which epistemic standard we accept.

Once Sally steps outside of her own standard to assert that S2 is rational for Susan but not for Sally, she is not in a position to claim that S2 is not maximally truth-conducive. Suppose that S2, for instance, generates more moderate confidences and fewer extreme ones than S1 given the same propositions and bodies of evidence. In the absence of any commitment one way or the other about whether it is better to have more or less moderate attitudes, S2 is clearly not superior to S1. Sally cannot assert the less than maximal truth-conduciveness of S2 without staying within the perspective of S1. However, as I have argued earlier, from within the perspective of S1, she cannot regard Susan's choice of S2 as anything but a mistake. Susan would be accepting S2 for what Sally regards as less than the best of reasons, namely on the basis of an imperfect scoring system. Outside of the perspective of S1, all Sally can say of S2 is that it makes a different set of trade-offs than S1, but is not obviously pareto inferior to it. She thus cannot say that S2 is not maximally truth-conducive. At most, she can say that it is differently truth-conducive without being more or less.

Once Sally denies that S1 is more truth-conducive than S2, she is vulnerable to the Arbitrariness Objection. If for her, S1 is not better than S2, she has no objection to changing epistemic standards by popping pills. However, forming beliefs by popping pills is arbitrary. Moreover, while Sally could defend accepting S1 on the basis that it matches the way she values acquiring truths and avoiding falsehoods while S2 does not, this precludes her from permissibly choosing S2. As I have mentioned earlier, I shall address this version of permissivism later in Chapter 6. Also, as I have mentioned earlier, if Sally believes that S1 is more truth-conducive than S2, she

³⁵ For instance, if the accuracy function is convex downwards (i.e. concave upwards), that can license believing both a proposition and its negation in a number of cases where the evidence is equivocal between the two.

cannot believe that it is rational for Susan to reason on the basis of S2. That is, Sally must accept uniqueness.

To be clear, all I have shown so far is that in accepting an epistemic standard, agents must accept uniqueness. I have not shown that uniqueness is in fact true. The distinction I am driving at was made by Kopec and Titelbaum (2016: 191) and is between acknowledged permissive cases and unacknowledged permissive cases. The thought here is that all the work in the Arbitrariness Objection and in my extension of it relies on the agent believing that a given piece of evidence is permissive. However, if the evidence is permissive but no agent believes that it is, then it would supposedly be possible for multiple agents to rationally adopt different standards, each mistakenly but rationally thinking that their own standard is the best and that everyone else is irrational. Then the Arbitrariness Objection would not apply since each agent accepts uniqueness and insists that her own standards are the correct ones. My only reply to such a move is that such a self-effacing version of permissivism commits us to being error theorists about rationality. I shall develop this point in greater detail in the next section (4.2).

If the above arguments are correct, the only way for permissivists to both reject uniqueness and avoid the Arbitrariness Objection is to endorse intrapersonal uniqueness about standards but suppose that there is some other factor that fixes a different epistemic standard as uniquely appropriate for each agent. Before I proceed to discussing such accounts of permissivism, I shall argue, in the next section, that the Evidence Pointing Problem can also be extended against versions of permissivism which are consistent with intrapersonal permissivism about standards.

4.2 Extending the Evidence Pointing Problem

The crux of the Evidence Pointing Problem argument is the claim that evidential support is univocal. The univocal-ness of evidential support can be presumed to be grounded in some fact. This fact might be a matter of certain empirical regularities that make it the case that the evidence E reliably indicates to some degree that a given proposition P is true (Goldman 2011; Ballantyne and Coffman 2011). In some cases, it might be a matter of logical entailment, probabilistic entailment³⁶ or some other *apriori* relation (Feldman 1995).

³⁶ See Ramsey (1931), Jeffreys (1948), Carnap (1950) and Howson (2003, 2007, 2009) for attempts at developing an account of logical probability. If logic can ground an account of

This point might require some explanation. Consider the following case which I call Empty Tank:

Empty Tank: The fuel gauge in my car indicates that there is no petrol in the tank. It has also been two weeks since I last filled up petrol. I conclude from my fuel gauge that the fuel tank in my car is empty.

If we were to ask a reliabilist why the fuel gauge reading supports the belief that the tank is empty, the reliabilist can point to causal connections between the tank and the gauge which make it the case that an empty reading on the gauge makes it more probable that the tank is empty. The causal connection between the tank and the gauge makes the gauge reading a reliable indicator of the amount of fuel left in the tank. This causal process therefore can be said to ground the evidential support relation on an indicator reliabilist picture. Given this causal relation, the gauge reading cannot support any other belief.

An evidentialist might deny that the fuel gauge reading alone supports the proposition that the tank is empty. Rather an evidentialist might say that the reading together with other evidence which indicates that the gauge tends to be accurate supports the proposition that the tank is empty.³⁷ For instance, it may be that every time the gauge reads empty, I need to fill in more than fifty litres of petrol before the tank is full. I may also know that the fuel capacity of my car is between fifty and fifty five litres. Moreover, I may observe that my fuel gauge decreases at a rate roughly proportional to the distance I drive and that I usually fill up with petrol every two weeks. All of these give me good reason to believe that my gauge is roughly accurate. This, together with the reading would then entail that the tank is probably empty. This entailment itself constitutes the evidential support relation. Given the entailment, the evidence cannot support any other belief.

To generalise, in making the claim that the evidence supports a given proposition, we presuppose that there must be some fact in virtue of which the evidence supports the proposition. Different accounts of justification will provide different accounts of what such a fact may be. The above example illustrates that however we conceive of such a fact, this fact ensures that the evidence cannot support any

probability or be extended to cover probabilistic cases, then there is something like an entailment relation that holds between the evidence and propositions in inductive and abductive cases.

³⁷ Conee and Feldman (2008: 90-91) argue that this way of conceiving evidential support prevents double counting.

other alternative. Let us call such facts that ground or constitute the evidential support relation G.

To illustrate, suppose, that in Empty Tank, the fact, which makes it the case that the fuel gauge reading supports the agent's belief that the tank is empty, is the causal relation between the amount of fuel in the tank and the reading in the gauge. Presumably, there is some device in the fuel tank which detects the fuel level and sends a corresponding signal to the gauge. Insofar as this device is working properly, the fuel gauge reading reliably indicates the level of fuel in the tank. In this case, G is the causal relation between the fuel level and the gauge reading. This causal relation makes it the case that the fuel gauge reading reliably indicates the fuel level. This is what, according to indicator reliabilism, makes the evidence support the belief that the fuel tank is empty. If indicator reliabilism is the correct account of justification, the evidence could not support any other doxastic attitude without the causal relation being different. For instance, the device in the fuel tank which detects the fuel level is broken and causes the gauge to read empty regardless of how much fuel is in the tank. However, the device cannot be both broken and working properly. The facts about what causal mechanism relates fuel levels to gauge readings fix whether and to what degree an empty gauge reading supports the proposition that the tank is empty. It follows that if the evidence does support the belief that the tank is empty, it cannot support any other doxastic attitude.

Another intuition that the Evidence Pointing Problem relies on is that the fact which grounds or constitutes the evidential support relation is not indexed to agents. Facts which are indexed to agents are things like the agent's epistemic standards, cognitive abilities, values, priors etc. Let me call such facts which ground or constitute the evidential support relation and which are not indexed to agents, F. In Empty Tank, the fact which grounded whether the gauge reading supported the belief that the tank is empty was the existence of the properly working device which established a causal mechanism between the fuel level and the gauge reading. This particular causal relation exists regardless of who happens to read the fuel gauge. Thus, in this case, the fact which grounds the evidential support relation is not indexed to persons. Therefore, if the evidence supports the belief that the tank is empty for one agent, Sally, that body of evidence supports that same doxastic attitude for all other agents who possess that evidence. Generalising from here yields uniqueness. If there is some fact that fixes what the evidence supports and

that fact is not indexed to persons³⁸, then there is no reason why different doxastic attitudes towards a proposition could be justified to people who have the same total body of evidence.

By contrast, suppose that at least some of the facts that grounded the evidential support relation were indexed to persons. Suppose, for instance, the question of whether a body of evidence supports a proposition depends on whether the agent can infer that proposition from the evidence via a reliable process. The question of whether a given process is reliable can depend on the cognitive capacities of the agent carrying out the process. To see why, consider the following situation.

Polling: Susan and Sally both possess a large amount of survey data regarding the upcoming election. Susan, who knows how to process the survey data can reliably infer that Candidate A will win. On the other hand, Sally, who does not know how to process the survey data cannot reliably infer that Candidate A will win. Therefore, she suspends judgment instead.

In Polling, there is a reliable process by which Susan can infer that Candidate A will win. However that process is not available to Sally. Therefore, if the facts that ground evidential support are at least partly constituted by the agent's cognitive capacities, then the evidence supports the proposition that Candidate A will win for Susan but not for Sally. It should be clear that even when the fact which grounds evidential support, G, is indexed to the agent, evidential support is univocal. If Susan has available to her a reliable process whereby she can infer that Candidate A will win and this fact makes it the case that believing that Candidate A will win is supported by her evidence, then suspending judgment or disbelief is not supported by the evidence for her. Likewise, the fact that no such reliable process is available to Sally would make it the case that, for Sally, her evidence supports suspending judgment about whether Candidate A will win.

Summing up, the existence of Gs, facts which ground the evidential support relation, is compatible with both uniqueness and permissivism. In fact, both permissivists and impermissivists are committed to the existence of such facts. To deny their existence is to suppose that there is no fact about the matter as to whether the

³⁸ Contra White (2005), Uniqueness presupposes nothing further than that the facts which ground or constitute the evidential support relation are not indexed to persons. While White presupposes that such facts are necessary, in principle, Uniqueness would be true even if the facts that ground evidential support are only contingent so long as they are not indexed to the agent.

evidence supports a particular doxastic attitude. Insofar as both permissivists and impermissivists are committed to the evidence actually supporting some doxastic attitude, there must be some fact which grounds this support relation. The core of the permissivist's disagreement with the impermissivist is whether all Gs are Fs; that is, whether any of the facts which constitute or ground the evidential support relation are indexed to persons. Uniqueness is true if and only if none of the facts which ground evidential support are indexed to persons.

As with the discussion regarding the Arbitrariness Objection, my targets in this section are accounts of permissivism like Schoenfield's which are consistent with intrapersonal permissivism about standards. I shall argue that such accounts cannot escape the Evidence Pointing Problem. In order to do this, we first need to settle what needs to be true of the facts which ground evidential support in order for intrapersonal permissivism about epistemic standards to be true.

Certainly, if all Gs are Fs, then intrapersonal uniqueness about standards must be true. Given that the evidence supports the same proposition for everyone, there is only one epistemic standard which would tell agents to believe according to the evidence. If only one standard is appropriate for everyone, then that standard is compulsory. However, even if some Gs were indexed to persons, this need not entail intrapersonal permissivism about standards.

Suppose that the facts were indexed to an agent's cognitive capacities. This would mean that facts about an agent's cognitive capacities would ground facts about what the evidence supported for her. If so, then this would fix for that agent which epistemic standard was appropriate. No other epistemic standard would be suited for that agent, given her cognitive capacities. Similarly, if the facts that grounded evidential support were indexed to an agent's values, then the agent's values would ground for that agent what the evidence supported. Then only one epistemic standard could tell her to believe what the evidence supported for her. While the agent's values might change at a later point in time, given the values she currently has, no other epistemic standard would be appropriate.

If the facts grounding or constituting evidential support are indexed to certain facts about agents, then against this background of facts, only one epistemic standard is appropriate for that agent. The only way for intrapersonal permissivism about standards to be true is if one of the facts which grounded evidential support was the agent's epistemic standards. If so, then while it would be trivially true that against the background fact of the agent's epistemic standard, only one epistemic standard

would be appropriate for her, the question of which epistemic standard an agent adopts is not a background fact.

In the ensuing, I shall argue that the facts which ground or constitute the evidential support relation cannot be indexed to the agent's epistemic standards. For ease of presentation, I shall centre my argument on Fs and concomitantly, uniqueness, thus ignoring for now the possibility that the facts which ground evidential support could be indexed to agents in one way or another. The argument can be easily extended to versions of permissivism consistent with intrapersonal uniqueness about standards. Ultimately, disregarding this possibility should not matter as I shall be showing in subsequent chapters that no such indexing proposal is plausible.

The first step of the argument involves showing that under Schoenfield's account, agents are committed to the existence of Fs. To see why, recall that for Schoenfield:

“[We] can just think of a set of standards as a function from bodies of evidence to doxastic states which the agent takes to be truth conducive. Roughly, this means that the agent has high confidence that forming opinions using her standards will result in her having high confidence in truths and low confidence in falsehoods.”
(Schoenfield 2014: 199)

This means that each agent must believe that her own epistemic standard is the most truth-conducive. After all, if she believes that there is some epistemic standard other than hers which would allow her to increase the rate at which she has high confidence in truths and low confidence in falsehoods but still sticks with her current standard, she would be irrational. If she believes her own epistemic standard is the most truth-conducive, then she ought to believe that the evidence really does support the conclusion that she would obtain by applying her epistemic standard.

To see why, note that an epistemic standard is a function from bodies of evidence to doxastic states. We can think of the epistemic standard as therefore consisting of a collection of rules of the form “If your evidence is E, take attitude A towards proposition P”. For the sake of simplicity I shall focus on just one such rule. In the following case, I shall suppose that there are seven instead of just three doxastic attitudes. Apart from suspending judgment, an agent may weakly believe or disbelieve, moderately believe or disbelieve or strongly believe or disbelieve in a proposition.

Suppose that an agent, Sally, accepts an epistemic standard S1 which includes the rule: if your evidence is E, moderately believe that P. Since she accepts S1, she believes that following the rule will be the most truth-conducive. It follows that she ought to think that E in fact moderately supports P. To see why this would be the case, suppose we were to deny that she ought to think that E moderately supports P and instead she believed that E weakly supports P. Sally would have the following combination of beliefs:

- A. Sally ought moderately believe that P in response to E
- B. Moderately believing that P is the most truth-conducive response to E
- C. E weakly supports P

A is just a restatement of the rule which is a part of S1. Believing that B follows from Schoenfield's account of what it means for Sally to accept S1, namely that she believes that her epistemic standards are the most truth-conducive. Once Sally accepts A and B, it is difficult to see how she could accept C. It is difficult to see how if E only weakly supports P, forming a weak belief that P would not be more truth-conducive than moderately believing that P. An explanation for why moderately believing that P is the most truth-conducive response is that E moderately supports P. Where there is a mismatch between claims about what the evidence supports and what is the most truth-conducive, there needs to be a further explanation apart from the evidence for why that response is the most truth-conducive. There also needs to be an account of what evidential support is such that weak support for P does not make weakly believing that P the most truth-conducive. However, it is unclear if any such account is available.

It should be noted that the above objection can be made with regards to any similar alternative to C. For instance, if Sally believed that E strongly supported P or that E supported not-P, there would still need to be some explanation as to why Sally did not think that strongly believing that P or disbelieving that P was the most truth-conducive. Therefore, Sally ought to believe that E moderately supports P. If Sally ought to believe that E moderately supports P, then she cannot permissibly believe that E supports P to any other degree. Therefore, according to Sally, moderately believing that P is the only rational response to E. If for each of her doxastic attitudes which are rationally held, she ought to believe that those attitudes are the only rational options, then she is committed to uniqueness.

To be clear, this not only commits Sally to uniqueness, it commits Schoenfield as well. To see why, suppose instead that all permissive cases were unacknowledged

permissive cases. That is to say, suppose that even though Sally, like all other agents, is committed to accepting uniqueness, uniqueness is in fact false. If uniqueness were false, then there would be no facts like F which would ground the evidential support relation and which would not be indexed to persons. The reason for this is that uniqueness is false only if there is no singular fact about what the evidence supports simpliciter. There is only what the evidence supports relative to people's epistemic standards, their capacities, values or priors. If there is no fact about what the evidence supports simpliciter, then there is no F. If there are no facts like F, but agents are committed to there being some facts like F, then they are mistaken about the kind of activity they are engaged in when they reason from evidence.

Being wrong about the existence of F is not merely being rationally committed to false beliefs. An agent's commitment to the existence of F goes to the core of what it means to make a rational inference on the basis of evidence. Recall that Sally thinks that moderately believing that P is the most truth-conducive response to E because E moderately supports P. Sally is committed to understanding rational inferences on the basis of evidence as being about believing according to what the evidence supports simpliciter. If there is no fact of the matter about what the evidence supports, then responding to evidence would be like astrology.

Astrologers, insofar as they are engaged in astrology, are committed to the existence of some fact which grounds an astronomical-psychological bridge principle. That is, they are committed to a principle that specifies how facts about the arrangements of celestial bodies in the night sky influence people's personalities. Insofar as there are no facts which could ground such a bridge principle, astrology is a nonsensical activity. By contrast, if there were such principles and hence facts grounding such principles, even if everyone tended to get these principles wrong, it would be merely difficult, but still not nonsensical. If reasoning on the basis of evidence commits us to the existence of F which does not exist, then reasoning on the basis of evidence is nonsense. However, surely since permissivists like Schoenfield would reject the claim that reasoning from evidence is a nonsensical activity, then they are committed to the existence of F and hence committed to uniqueness.

If we wish to deny that accepting an epistemic standard commits an agent to F, we need to say that it commits them to facts which are indexed to agents instead. Given that Schoenfield's account of the evidential support relativizes evidential

support to people's epistemic standards, we might try applying this account to such facts. That is to say, suppose that instead of C, Sally believes C* instead.

C*. E moderately supports P relative to S1

However, this move makes Schoenfield's account circular. C* is intended to explain A and B. To see why, recall that A is a rule that is extracted from S1 and S1 explains why E supports moderately believing that P for Sally but not for Susan. B explains A in the following sense. If Sally were to be asked why she should moderately believe that P in response to E, she cites her belief that doing so is the most truth-conducive option. However, when asked why she thinks moderately believing that P is the most truth-conducive option, she cannot cite C*. If she were to cite C* she would be using C* to explain why C* is the case. If she does so, she would be explaining why a standard recommends the attitude it does by appealing to the fact that the evidence supports that attitude relative to the standard. But by her own lights, the fact that one's own standard recommends a given response to the evidence explains why the evidence supports it, relative to the standard. Thus C* would be used to explain C* and that cannot be right.

Whatever the fact that the truth-conduciveness of an epistemic standard is to be grounded in is, that fact cannot be indexed to the standard. We cannot explain why a given epistemic standard is truth-conducive by appealing to that epistemic standard. Therefore it has to be grounded either in facts which are not indexed to persons or to facts like the agent's values, priors or cognitive capacities. In the next three chapters, I shall discuss each of the latter three indexing proposals and argue that each of them face certain distinct problems. For now, all I need to note is that if the fact that grounds the truth-conduciveness of the epistemic standard is one like the agent's values, priors or cognitive capacities, intrapersonal permissivism about standards cannot be true. As can be seen in the Polling example, if facts about the agent's cognitive capacities ground the degree to which her evidence supports a proposition, then only the epistemic standard which requires her to form the corresponding doxastic attitude is rationally permissible for her.

In this chapter, I have argued that the Arbitrariness Objection and the Evidence Pointing Problem can be extended to apply to versions of permissivism in which more than one epistemic standard is permissible for an agent. As mentioned in the previous chapter, a central task of this thesis is to show that for each burden of judgment which presupposes permissivism, the disagreement caused by that burden is not rational. I argued that the second burden, which essentially claimed

that people could rationally disagree because they had different epistemic standards, was crucial in making rational disagreement possible because of the way in which having different epistemic standards solved the Arbitrariness Objection and the Evidence Pointing Problem. In showing that these two objections can be extended to versions of permissivism consistent with intrapersonal permissivism about epistemic standards, I have closed off one possible route by which the second burden of judgment could generate rational disagreement. That is to say, people cannot permissibly disagree about the strength of an evidential consideration simply because more than one epistemic standard is rationally permissible for a given agent. If it is ever permissible for people to disagree about the strength of their evidence, this could only be the case if some other burden of judgment made different epistemic standards uniquely appropriate for different agents. The only versions of permissivism left are ones which are consistent with intrapersonal uniqueness about standards: For any given agent, only one, but not necessarily the same epistemic standard is permissible for her.

Chapter 5: Permissivism, Cognitive Capacities and Limitations

Background

In Chapter 3, I argued that intrapersonal uniqueness is true and that the most plausible version of permissivism was one according to which the doxastic attitude which was rational for a person was relative to the epistemic standard they adopted, provided that the epistemic standard was itself permissible for the agent. In Chapter 4, I argued that versions of permissivism which are consistent with intrapersonal permissivism about epistemic standards are also vulnerable to the Arbitrariness Objection and the Evidence Pointing Problem. The only views left are those which are consistent with intrapersonal uniqueness about epistemic standards: for each agent there is only one rationally permissible epistemic standard, but not necessarily the same one. This raises the question as to why different epistemic standards would be uniquely rational for different agents. Any satisfactory answer to this must provide some respect in which agents can differ from one another, and some account of how this aspect in which they differ makes a difference to which epistemic standard they should employ. I shall cover three such proposals: The first of which I shall cover in this chapter is that the question of which epistemic standard is most truth-conducive for an agent depends on her cognitive capacities. The second proposal which I shall cover in the next chapter is that how we measure the truth-conduciveness of an epistemic standard depends on how we value the twin epistemic goals of pursuing truths and avoiding falsehoods. The third proposal, which I shall cover in Chapter 7, is that the epistemic standard we ought to employ depends on what doxastic attitudes we have prior to acquiring the evidence.

5.1 Truth Conduciveness Relative to Cognitive Capacities

As it is impossible to cover every possible variation in the space I have available, I shall cover what I take to be versions representative of the most dominant theories in epistemology. One version of permissivism says that people may adopt different epistemic standards because they have different cognitive abilities. The existence of these different cognitive abilities means that different epistemic standards will be maximally truth-conducive for different people. In this chapter, I shall be addressing Simpson's account, which depends on people being cognitively imperfect in one way or other.

Simpson illustrates his account of permissivism with a rather lengthy case he refers to as Detective. I shall reproduce the case in full.

Detective: “Veronica Mars and Nancy Drew are investigating a crime for which X and Y are the prime suspects. After thoroughly reviewing the same body of relevant evidence, Mars believes it’s 70% likely X is the culprit and 30% likely Y is the culprit, whereas Drew believes it’s 30% likely X is the culprit and 70% likely Y is the culprit... Suppose firstly that Mars and Drew employ different epistemic standards, i.e. they use different methods for assessing and interpreting their evidence in order to form their beliefs. Suppose, for instance, that they have different methods for evaluating the probative force of testimonial evidence, a type of evidence that’s often an important part of the data they examine as detectives. Mars is excellent at recognizing insincere testimony – she’s very sensitive to various subtle verbal and non-verbal tell-tale markers of insincerity – and thus she interprets different instances of testimony differently, ascribing a different probative force to insincere testimonial evidence. Drew is not very good at identifying insincere testimony, but she is excellent at making holistic, gestalt assessments of the combined probative implications of large bodies of disparate evidence. So unlike Mars, Drew doesn’t pre-sort the testimonial evidence, based on its sincerity. Where Mars reviews the testimonial evidence in sequence, identifying each item’s probative force in its own right, Drew reviews all testimonial evidence in the same light, and assesses the probative force of the testimonial data as a set. Mars applies her standards because she’s cognitively well-equipped to apply them. She has the kind of perceptual and attentional abilities that enable her to accurately judge the sincerity of testimony. In a similar way, Drew is cognitively well-equipped to apply her standards. She has a particular imaginative facility that enables her to formulate credible explanatory narratives, based on gestalt interpretations of diverse bodies of evidence... Although Mars and Drew are working together for the first time, they have a long track-record of working independently on other cases. And in light of their track-records, both of them have good reason to believe that when they apply their own epistemic standards, in forming beliefs about cases relevantly similar to the case at hand, they are very likely to get the answer right. In short, each of them is in the position of (i) having a combination of cognitive abilities and epistemic standards which reliably generates

accurate beliefs, and (ii) having good reason to believe that this is the case.” (Simpson 2017)

Simpson argues that this case provides an example of how people with different but equally reliable cognitive capacities could adopt different epistemic standards. If Simpson is correct, people who have different cognitive capacities could adopt different standards and therefore have different responses to the evidence which were nevertheless still rational. Simpson’s account relies on the presumption that the epistemic standard best suited for each person is maximally truth-conducive for that person. If Nancy Drew were to try to apply Veronica Mars’ standards she would do less well because she would do badly at evaluating the sincerity of the testimony. Similarly, Veronica Mars would also perform sub-optimally if she applied Nancy Drew’s standards, as she is prone to make mistakes when trying to assess the probative force of the evidence as a whole. Both standards play to their respective agents’ strengths and mitigate their weaknesses.

However, it is this aspect of the situation which undermines the rationality of both detectives’ responses. Both Nancy Drew and Veronica Mars fail to respond perfectly well to all their evidence because both have weaknesses: parts of the evidence which they are unable to reliably assess well. To see why, consider an amalgamation of Nancy Drew and Veronica Mars: Nancy Mars. Nancy Mars has both Nancy Drew’s and Veronica Mars’s weaknesses: she is unable to reliably assess the sincerity of testimony and evaluate the evidence as a whole. Nancy Mars is clearly less reliable than both Nancy Drew and Veronica Mars. Nancy Mars is less reliable than Veronica Mars due to the reason that she is unable to reliably assess the sincerity of testimonial evidence; hence, she tends to overestimate the probative force of testimony from insincere people and underestimate the force of testimony from those who are sincere. Since Nancy Mars differs from Veronica Mars in just this one respect, the former is clearly less reliable than the latter. Plausibly, to over or underestimate the strength of a given piece of evidence is, by definition, to make a mistake in reasoning. If reasoning well necessarily involves forming beliefs in proportion to the strength of the evidence, then forming beliefs that deviate from what the evidence points to is going to be a mistake. As I shall be arguing in the following sections, if a belief can only be formed on the basis of a mistake, then it is not rational given the evidence.

Since Veronica Mars is just as reliable as Nancy Drew, Nancy Mars must be less reliable than Nancy Drew. The only difference between Nancy Mars and Nancy

Drew is that the latter evaluates her evidence as a whole, whereas the former evaluates her evidence piece by piece. Since evaluating the evidence piece by piece is less reliable than evaluating it as a whole, evaluating it piece by piece must involve in at least some cases mis-estimating the strength of the evidence as a whole. The assumption here is that to evaluate one's evidence properly is to evaluate it in a maximally truth-conducive manner. I shall evaluate this assumption later in this chapter, but, for now, if it is true, then Nancy Mars must make at least some mistakes in reasoning by evaluating her evidence piece by piece rather than as a whole.

If evaluating one's evidence piece by piece results in a mistake in a given case, then Veronica Mars would make a mistake in reasoning in that case. Similarly, if failing to account for the sincerity of testimony is a mistake, then Nancy Drew also makes a mistake whenever she applies her epistemic standards. As such, since they both make mistakes in reasoning, their disagreement cannot be rational. Notably, they must be substantive mistakes because they involve misestimating the overall strength of the evidence. The doxastic attitude they end up with arguably does not fit the evidence. If, on the other hand, they were to properly evaluate the evidence by both reliably assessing the sincerity of testimony and evaluating the evidence as a whole, there seems to be no particular reason why they would disagree.

Resisting this argument requires showing that evaluating the evidence piece by piece and failing to assess the sincerity of testimony are not mistakes in reasoning; or even if they are mistakes, do not render the resultant belief irrational. However, if they are not mistakes then it would be permissible for Nancy Drew and Veronica Mars to adopt Nancy Mars's epistemic standard. Yet if more than one epistemic standard was permissible for any agent, it would, as discussed in Chapter 4, be arbitrary as to why she uses one standard rather than another. Moreover, as discussed in the same chapter, the facts which ground evidential support cannot include the agent's epistemic standard. If so, more than one epistemic standard cannot be permissible for any given agent. What is implicit in Simpson's argument is the claim that Nancy Drew can permissibly not assess the sincerity of testimony because she is not able to reliably do so. Similarly, Veronica Mars can permissibly assess the evidence piece by piece because she is incapable of assessing it as a whole.

What is implicit in the Detective case is the idea that the reason why people can have different epistemic standards when they have different cognitive capacities is that they can permissibly lower their epistemic standards in response to their own cognitive limitations. Thus, since Nancy Drew cannot reliably assess the sincerity of testimony, she ought to adopt an epistemic standard according to which she should ignore all evidence pertaining to sincerity. If, on the other hand, people's epistemic standards need not be lowered in response to cognitive limitations, there would be no reason to think that a standard according to which Nancy Drew need not assess the sincerity of testimony is appropriate for her just because she cannot reliably assess the sincerity of testimony. Likewise, there would be no reason to think that a standard according to which the agent need not evaluate the evidence as a whole is permissible for Veronica Mars.

5.2 Difficult Questions and Cognitive Limitations

Simpson's account of how different cognitive capacities make appropriate different epistemic standards ultimately boils down to a variation of the claim that people can rationally disagree because they have limited cognitive resources or capacities. This claim is made explicitly by Gaus (2011) and Rosa (2012; 2016), and is also implicit in Rawls's (1993) account of the first, third and fifth burdens of judgment.

Recall that according the first burden of judgment, people can rationally disagree because "the evidence... bearing on the case is conflicting and complex and thus hard to assess and evaluate" (Rawls 1993: 56). Similarly, according to the fifth burden, "there are different kinds of normative considerations of different force on both sides of an issue and it is difficult to make an overall assessment" (ibid: 57). As I argued in Chapter 2, both the first, third and fifth burdens claim that because it is difficult to evaluate the overall normative force of one's evidential or moral reasons, disagreement about these issues, even given the same total body of evidence, can be rational.

My claim, in this section, is that disagreement that comes about because the overall normative force of our reasons is difficult to assess would be rational only if agents should lower their epistemic standards because of their cognitive limitations. While Rawls himself says little on this topic, Gaus's (2011) account of reasonable disagreement provides a plausible account linking the difficulty of an issue to the rationality of disagreement about that issue. Gaus's argument consists of three claims.

- A) People have limited cognitive resources.
- B) Avoiding errors in reasoning can in some cases require large amounts of cognitive resources.
- C) There is a point beyond which people are not obligated to use their cognitive resources to avoid errors in reasoning.

These three claims together imply that when matters are difficult to assess, disagreement can be reasonable. In order to show why this is the case, I shall elaborate on each claim in turn.

Gaus's first claim is that people have limited cognitive resources. Reasoning requires time and effort and people only have a finite amount of time and effort that they can spend on reasoning. People cannot continue reasoning indefinitely; they eventually have to act on the beliefs that they have reasoned about so far (ibid: 246-247). Gaus's account is concerned with how much reasoning an agent does, that is, how much of an agent's system of beliefs or evidence she looks through to identify supporting and defeating considerations. Apart from identifying a consideration, we might suppose that cognitive resources are also expended in assessing the strength of that consideration. It takes at least some effort for an agent to make herself aware of her biases and screen them off when assessing how significant a given piece of evidence is. We may also think that correspondingly that the amount of effort a person can continuously exert over a given period of time is limited as well.

In addition to the claim that cognitive resources are limited, Gaus's second claim is that avoiding errors in reasoning can consume large amounts of cognitive resources. Consider that when we casually examine a piece of evidence, we are prone to hastily over or underestimate the strength of that consideration. For instance, as Kornblith notes, we are often prone to drawing conclusions about a population on the basis of a very limited sample size. He points to the common tendency to draw conclusions about consumer products on the basis of a single person's bad experience (Kornblith 1983: 42-44)³⁹. This tendency seems to indicate that refraining from hastily forming beliefs on inadequate inductive evidence takes more effort than many among us are willing to expend. Gaus notes that the amount of reasoning one has to engage in so as to take into account all the relevant considerations found in an agent's system of beliefs exceeds the capabilities of ordinary reasoners. He says that "only ideally rational agents, capable of scanning

³⁹ To be sure, Kornblith (1983) does try to argue that such inferences are justified in situations where we cannot obtain better evidence. However, he notes that these inferences still run afoul of what he calls rules of ideal reasoning.

their entire belief system for lurking defeaters” could completely avoid errors in reasoning (Gaus 2011: 146). We need not agree with Gaus in thinking that every question will require us to scan through our entire system of beliefs in order to avoid errors in reasoning. Perhaps in particularly simpler cases, it is fairly obvious what the evidence regarding the proposition can be, and scanning through this set of beliefs is within the grasp of most ordinary agents. For instance, splitting the dinner bill equally between two persons is something that most ordinary agents are capable of and does not require anyone to scan through their entire system of beliefs. Other questions might require more reasoning and this would mean that the agent had to scan more of her system of beliefs to identify the relevant considerations. Summing up, reasoning well requires cognitive resources, and cases which require more reasoning or more careful reasoning require more cognitive resources. Given that people have limited cognitive resources, questions which use up most of if not more cognitive resources than a person may have available will be difficult to assess without committing some kind of error in reasoning. Therefore, the normative force of an agent's reasons would be difficult to assess only because of her cognitive limitations.

If the difficulty of assessing the normative force of one's reasons can render the resulting disagreement rational, this can only be because people, due to their cognitive limitations, ought to accept epistemic standards which are constituted by relatively less demanding epistemic norms. If agents could not permissibly lower their standards in this way, then, they ought to reason according to the more demanding epistemic norms even though they are cognitively limited. If so, the fact that they are cognitively limited and the reasons are therefore difficult to assess would play no role in making any resulting disagreement rational. Therefore, in order for the first and fifth burdens to be sources of rational disagreement, it must be permissible for agents to employ less demanding epistemic standards in response to their cognitive limitations.

This must similarly be the case for the third burden of judgment, according to which people rationally disagree because their concepts to some degree or another are vague and subject to borderline cases. Williamson explains the existence of vague predicates by appealing to cognitive limitations that prevent us from clearly identifying whether a given borderline case falls under the predicate. As I suggested in Chapter 2, these cognitive limitations are of the same sort that could make inferences when the evidence is conflicting and complex subject to rational disagreement. If so, it must also be the case that in order for the third burden of

judgment to generate rational disagreement, it must be permissible for agents to employ less demanding epistemic standards in response to their cognitive limitations.

As mentioned earlier, cognitive limitations also play a crucial role in Simpson's account. Indeed in reference to Detective, Nancy Drew is unable to reliably assess the sincerity of testimony and Veronica Mars similarly cannot assess the evidence as a whole, and has to proceed piece by piece. Their disagreement is rational only if it is permissible for each to employ epistemic standards which are less demanding in ways that correspond to those respects in which they are cognitively limited.

5.3 Limitation of Epistemic Obligations by Cognitive Capacity

It is precisely this last claim which is defended by Gaus. According to the third claim made by Gaus, there is some point beyond which people are not obligated to use their cognitive resources to avoid errors in reasoning. Gaus, for instance, endorses the following account of when individuals have sufficient reason to hold a particular doxastic attitude. Let's call this principle Provisionally Sufficient Reasons:

Provisionally Sufficient Reasons: "Alf has (provisionally) a sufficient reason R^{40} if and only if a 'respectable amount' of good reasoning by Alf would conclude that R is an undefeated reason" (ibid: 250)

Provisionally Sufficient Reasons implies that people do not need to engage in a more than respectable amount of reasoning in order to rationally hold a particular doxastic attitude towards a proposition. After all, if a given, "respectable" amount of reasoning is sufficient for the resulting doxastic attitude to be rational, then a more than sufficient amount of reasoning cannot be obligatory. We might, in addition, think that apart from the amount of reasoning, there could also be some "respectable" amount of effort we put into reasoning beyond which there is no further obligation to ensure that we do not under or over-estimate the strength of a given piece of evidence.

There are different ways in which we could determine what level of cognitive resources we have to commit in order to have engaged in a respectable amount and quality of reasoning. For instance, it could be that people have an obligation to use all their available cognitive resources. If people have different amounts of cognitive resources available, what counts as a respectable amount of reasoning

⁴⁰ In case it is not already clear, I will take the elocutions "A has sufficient reason to believe P" and "It would be rational for A to believe P" to be interchangeable.

would differ between them. Or, it could be that people ought to expend some fixed fraction of their total cognitive resources in reasoning, but no obligation to expend more than that. Alternatively, there could be some fixed amount of resources that people are obliged to expend. People who have very little in the way of cognitive resources would be obliged to expend all or most of their cognitive resources while those who are much wealthier in terms of cognitive resources need to spend only a small fraction of their resources before they have satisfied their epistemic obligations. I shall discuss these different accounts later in this chapter. Suffice it to say that for any of these accounts, it will sometimes be the case that a more than the obligatory amount of cognitive resources needs to be expended in order to avoid errors in reasoning.

Suppose that in a given case, a more than obligatory amount of resources is necessary in order to avoid erroneous reasoning. Alfred and Betty share a body of evidence which contain two pieces of evidence, R1 and R2 which pertain to whether P. Suppose that after a “respectable” amount of reasoning, Alfred uncovers within his evidence R1 which supports a belief in P, but not R2. Betty, who shares the same body of evidence uncovers after a “respectable” amount of reasoning R2 which supports disbelief in P. Since the amount of reasoning required to account for all the considerations, R1 and R2, relevant to P exceeds the amount of reasoning that Alfred and Betty are obligated to put in when considering whether P, we can treat this as being a case in which the question of whether P is difficult to assess. Moreover, since Alfred and Betty have both put in a respectable amount of reasoning, according to Provisionally Sufficient Reasons, they would have sufficient reason and are therefore rational in their corresponding beliefs. This example should make clear at least one way in which a proposition being difficult to assess creates room for rational disagreement. We can generate other examples whereby the amount of cognitive resources required to completely avoid errors in reasoning exceeds the amount which agents are obligated to expend in reasoning. In any such case, the fact that the cognitive cost of completely avoiding errors is so high, combined with the supposed fact that people may permissibly expend fewer resources, makes it the case that rational disagreement is possible.

It is difficult to dispute that the cognitive costs in completely avoiding errors in reasoning can be very high in some cases. The more questionable assumption in a view like Gaus’s is that it is sometimes epistemically permissible to make errors while reasoning about a proposition. In order to argue that difficulty in assessing whether a proposition is true does not give rise to rational disagreement, I submit

that making an error in reasoning is always irrational. In addition, I shall argue that it is always an error to believe that P, when I have at least as strong a reason not to. Let us call the combination of these claims the Naïve account of epistemic rationality. We can contrast the Naïve account with what I call a Forgiving account according to which there are some cases in which I rationally hold a given doxastic attitude even if I make an error in reasoning or there are some cases where I do not make an error even though I believe that P even though I have at least as strong a reason not to. I shall first sketch out an argument for the Naïve account, followed by some objections to it. I shall subsequently reply to these objections to show that the Naïve account is in fact true.

5.4 The Naïve Account

Intuitively, if I make a mistake in reasoning my way to a particular belief, I have formed that belief in a rationally defective way. If I have formed a belief P in a rationally defective way and nothing further has changed, then I am not rationally believing that P. That is to say, my belief that P is not doxastically justified. According to Turri (2010: 315), doxastic justification pertains to an evaluation of the performance of an agent. On the more traditional view of doxastic justification, a belief is doxastically justified only if it has been formed on the basis of evidence which supports it (Swain 1979; Pollock and Cruz; 1999; Korcz 2000; Feldman 2002; Kvanvig 2003). Here too, doxastic justification is an evaluation of the process by which the belief is formed. Either way, insofar as there are performance errors, then we cannot evaluate the performance as being without defect. That is to say, if there are performance errors in the process of acquiring a belief, that belief is not doxastically justified.

A quick objection to this argument is that while it is true that a performance error entails that the process by which the belief was formed was defective in some way, the question we are interested in is another sense of justification which has more to do with whether the belief is consistent with or fits the evidence. That is to say, permissivism is not a claim about whether two people who disagree are both doxastically justified in their beliefs, it is a claim about propositional justification. However, this objection is ultimately unsuccessful. To see why, I shall first briefly survey what I take to be the dominant accounts of propositional justification.

The traditional view of propositional justification is that a proposition, P is propositionally justified for an agent, Alice, if and only if she has sufficient epistemic reasons R to believe P (Swain 1979; Pollock and Cruz 1999; Korcz 2000). Turri

(2010) presents a series of cases where intuitively, a belief which has been formed on the basis of one's evidence is not doxastically justified. In one case, a juror, Improper, looks at her tea leaves and concludes that the tea leaves say that the evidence supports a guilty verdict. In fact, the evidence does in fact support a guilty verdict. In another case, an agent, Lacy, could use *modus ponens* to infer that R from his evidence, but instead infers that R by using a different invalid inference rule. These cases, according to Turri, indicate that doxastic justification is more than just a matter of believing on the basis of evidence, but is irreducibly about reasoning well from the evidence (ibid). Turri then defines propositional justification in terms of doxastic justification: P is propositionally justified for Alice if and only if there were some way of reasoning to a belief that P which is available to her and which if she took, she would be doxastically justified in believing that P (ibid).

While there is considerable disagreement about the relation between doxastic and propositional justification, everyone agrees that if an agent is doxastically justified in believing that P, P is propositionally justified for her. This is because on the traditional view, if a belief is doxastically justified, it has to be propositionally justified and based on the evidence in virtue of which it is propositionally justified. On Turri's view, a belief that P is propositionally justified for an agent if and only if she has available to her some path of reasoning which would result in her having a doxastically justified belief that P. When she actually has a doxastically justified belief that P, she clearly has a path of reasoning available to her which would result in her being doxastically justified in believing that P. Either way, doxastic justification entails propositional justification. It follows that if P is not propositionally justified for an agent, then her belief that P is not doxastically justified. To see why the Naïve account is true, let me introduce an initially plausible account of propositional justification which I shall call the Simple view:

Simple: A proposition P is propositionally justified for an agent Alice if and only if her total body of evidence supports P.

Simple seems to cover the main instances of when we think propositions are or are not justified for someone. Typically, when we suppose that a proposition is justified for someone, we think that this is so because it is supported by their evidence. And similarly, where it is not justified for someone, we think it is not because it is not supported by their evidence. Consider again the case where Alfred and Betty share a total body of evidence consisting of R1 and R2 where R1 supports P and R2 supports not-P respectively. Suppose R2 is stronger than R1. The evidence as a

whole therefore supports not-P. Given that Alfred's evidence consists of R1 and R2, P is not propositionally justified for him while not-P is. Therefore, believing that P cannot be doxastically justified for him. Simple therefore seems to imply that Alfred's failure to take into account R2 is an error in reasoning.

As noted earlier, the first, third and fifth burdens cause disagreement when a matter is difficult to assess. When matters are difficult to assess, disagreement arises because limited cognitive resources make it such that agents fail to consider some of their reasons or because they mis-estimate the strength of their reasons. When either of these occur, the resultant belief will not match what the evidence supports except by coincidence. For instance, normally if I were to overestimate the strength of one piece of evidence, my overall assessment of the strength of the evidence will likewise be skewed. It is of course possible for one mis-estimation to perfectly cancel out another, but we cannot count on such lucky accidents occurring. Other than these coincidental cases, any belief subject to the first, third and fifth burdens will not be propositionally justified. Therefore they cannot be doxastically justified either.

The only explanation for this is that there must have been some error in reasoning which led to the agent having a belief that was not supported by the evidence. This error in reasoning would be some inferential step that was taken or omitted which led to a conclusion that differed⁴¹ from what is supported by the evidence as a whole. For instance, by failing to take into account R2 which was part of his evidence, Alfred believed that P instead of not-P. This failure to account for R2 is therefore a reasoning error on Alfred's part and this error explains why his belief that P is not doxastically justified.

To sum up, on the Naïve view, the first, third and fifth burdens generate errors in reasoning. These errors in reasoning suffice to make it the case the resultant belief is not doxastically justified. These errors in reasoning also cause the resultant belief to differ from what would be supported by the evidence. The reason for this is that on the first, third and fifth burdens, cognitive limitations can cause agents to fail to consider all of their evidence or to mis-estimate the strength of their evidence. Since the resultant belief is not supported by the evidence, as per Simple, it is not

⁴¹ I am not committing to whether any similar inferential step which would not make a difference counts as an error. I leave open both the possibility that Betty makes a mistake by failing to take into account R₁ and the possibility that she does not just because R₁ is weaker than R₂.

propositionally justified either. This means that the first, third and fifth burdens cannot be sources of rational disagreement.

5.5 Forgiving Accounts

Arguing against the Naïve account will necessarily involve denying Simple. So long as Simple is true, any inferential step which is taken or omitted due to cognitive limitations is going to count as an error in reasoning and will not result in a propositionally justified belief. For instance, Rosa (2016) defends permissivism on the basis that propositional justification does not amount to “sheer evidential support” (ibid: 96). Rosa takes propositional justification of P to be a matter of whether an agent could be doxastically justified in believing that P. He explicitly references Turri’s account of propositional justification according to which:

Propositional Justification (PJ): “Necessarily, for all S, p, and t, if p is propositionally justified for S at t, then p is propositionally justified for S at t because S currently possesses at least one means of coming to believe p such that, were S to believe p in one of those ways, S’s belief would thereby be doxastically justified.” (Turri 2010: 320)

According to PJ, propositional justification of P for an agent is dependent on whether she could in fact come to be doxastically justified in believing that P. For Rosa, the question of whether P is propositionally justified for someone depends on whether she is capable of becoming doxastically justified in believing that P. That is to say, Rosa endorses a “may” implies “can” principle⁴² for epistemic justification. If cognitive limitations preclude an agent from believing that P without making a mistake along the way, P would not be propositionally justified for her even if her evidence did in fact support P. On Rosa’s account, each agent’s cognitive capacities set the threshold for how demanding the requirements of rationality are. An agent is obligated at most to expend those cognitive resources she has in reasoning from her evidence, not more. Kornblith (1983; 2001) endorses a similar view whereby what an agent ought to believe in light of the evidence is informed by her cognitive capacities. The thought here is that it should not be impossible for us to infer in those ways we are epistemically entitled to.

⁴² Actually Rosa calls it a “can” implies “can” principle (2016: 97) where the first “can” refers to permission and the second “can” refers to capacity. My formulation of it removes the ambiguity about the word can.

One consequence of such a view is supposed to be that there are a number of cases in which our evidence supports a proposition, but we are permitted to suspend judgment. Rosa (2016) points out that for any highly complicated mathematical theorem that is beyond our abilities, there is no performance error involved in failing to evaluate whether it is true or false, even if we already justifiably accept the axioms and the theorem follows logically from those axioms. For instance, everyone's evidence, according to Kornblith, supports the proposition that mathematics is incomplete (Kornblith 2001: 233-234). However, it took Gödel to establish this, and intuitively most who have not heard of his proof may permissibly suspend judgment about it since we are not capable of working out that mathematics is incomplete. Kornblith also points out (*ibid*) that this is not limited just to esoteric mathematical theorems, but even to empirical questions where only a genius could recognise the connection between the evidence and the supported proposition. Intuitively, if we are not capable of making the inference, we ought to suspend judgment. That is to say, suspending judgment seems doxastically justified in these cases. More strongly, it seems obligatory for those of us unable to recognise these evidential connections. However, if we ought to suspend judgment, then P cannot be propositionally justified to us even if it is supported by our evidence. Therefore, or so the argument goes, Simple must be false.

However, this view, according to which less demanding epistemic norms apply to agents who are more cognitively limited, is not without theoretical costs. One significant cost is that people with very limited cognitive abilities would get away with very bad reasoning. That is to say, the epistemic standards appropriate for someone with very limited cognitive abilities would not resemble, to any significant degree, any plausible picture of what good reasoning looks like. Turri (2010) brings up the example of Ron, who is psychologically incapable of believing that the US should not invade Iran. He knows that invading Iran would be catastrophic and knows that if it is catastrophic, the US should not do it. It intuitively seems that the proposition that the US should not invade Iran is propositionally justified for Ron. However, since Ron is psychologically incapable of believing that the US should not invade Iran, he lacks the means to become doxastically justified in believing it. According to PJ, this seems to mean, that believing that the US should not invade is not propositionally justified for Ron. Depending on exactly how extensive Ron's psychological incapacities are, Ron could be propositionally justified in believing all sorts of things so long as his reasoning capacities are severely bad. For instance, if

he were incapable of applying *modus ponens*, or the disjunctive syllogism, there are many beliefs Ron may justifiably have despite glaring failures to reason adequately.

It is by no means the case that no one is unwilling to bite this bullet. Rosa (2012; 2016) for instance presents the following case as one in which an agent can rationally suspend judgment.

Virus: “Suppose Amanda justifiably believes both:

- (i) If I clicked the wrong link, my e-mail has a virus now.
- (ii) My email has no virus now.” (Rosa 2012: 574-575)

According to Rosa, Amanda, who is unable to carry out *modus tollens*, is justified in suspending judgment about whether she clicked the wrong link.

Nevertheless, this still seems implausible as both Amanda and Ron appear severely intellectually disabled. It is a banal observation that most of us do not reason to the absolute best of our ability. If what is justified for us depends on our cognitive capacities, then those whose cognitive capacities are the most limited are the most rational. These would be people whose cognitive capacities are so limited that they are currently already reasoning to the best of their capacity; they do not have room to underperform. However, it is absurd to think that the reasoning of those who are least able to reason is rated the best among all of us. Yet, such a cost is unavoidable, if we peg the demandingness of an agent’s epistemic standards to her cognitive capacities.

One way to get around this objection is to peg our epistemic standards to people’s average cognitive capacities. Then, we are still able to say of the average person, that she may suspend judgment when she encounters complicated cases like the Gödel’s incompleteness theorem because the average person is not able to work out that Gödel’s incompleteness theorem is true. On the other hand, we can say of people like Ron and Amanda, that they are not justified in suspending judgment in their relatively simpler cases, because the average person, given the same evidence, is capable of working out that the US should not invade Iran and that Amanda did not click on the link.

One significant objection to the second proposal is that even though pegging epistemic standards to the average capacity of agents manages to avoid some counterintuitive results, there are still some cases which it will find difficult to handle. The key case here involves people with far higher than average cognitive ability

(Turri 2010). Consider what the average ability account says about what is justified for someone like Gödel. As a genius, Gödel was able to work out that the incompleteness theorem was true. However, the average person is not. If we peg epistemic standards to the average human cognitive capacity, then the theorem is not justifiable to Gödel, even though it clearly has to be. This is because, as per the standard pegged to the average cognitive capacity, Gödel ought to suspend judgment about the incompleteness theorem. However, this is absurd! If Gödel is indeed capable of soundly reasoning his way to seeing that the theorem is true, it cannot be impermissible for him to believe it when he has done the reasoning correctly.

Perhaps the problem lies instead with the claim that Gödel ought to suspend judgment. Perhaps the average-capacity-standard only permits him to suspend judgment. It could also permit him to believe the theorem is he has done the reasoning correctly. However, this violates intrapersonal uniqueness. As discussed in Chapter 3, the solution to White's Arbitrariness Objection was to claim that only one attitude was justified given the evidence relative to a particular epistemic standard. By relativizing justification to an epistemic standard, the agent finds it inappropriate to acquire beliefs by flipping a coin or popping a pill. However, if even relative to the average-capacity-standard, Gödel was permitted to either suspend judgment or believe the incompleteness theorem, then there is no reason for him to prefer reasoning on his evidence to acquiring beliefs by randomly popping a pill. Therefore, more than one attitude towards the theorem cannot be permitted for Gödel given the average-capacity-standard. Given that standard, Gödel ought to suspend judgment and as we discussed in the previous paragraph, this is absurd.

Turri (*ibid*) solves this problem by adopting a heterogeneous account. For people of average and below average ability, what they ought to believe is limited to what people on average are capable of. For those with above average cognitive capacities, what they ought to believe, given the evidence, is limited only by their own cognitive capacities, accommodating therefore, three sets of intuitions. The first is the intuition that people of normal cognitive ability may permissibly suspend judgment in complicated cases. The second intuition is that people of below average cognitive ability may not permissibly suspend judgment in relatively simpler cases even when those are still beyond their personal ability to reason through. The third intuition is that people with extremely high cognitive ability would be justified in believing what is supported by their evidence even in complicated cases.

Even if this accommodates all three intuitions about cases, this third heterogeneous account is not without its set of problems. Note firstly that both the second and third accounts lack adequate justification. Recall that limiting an agent's obligations to her own cognitive capacities was justified by a plausible principle: "ought" and "may" imply "can". The principle is grounded in the intuition that it is incoherent to demand of someone, something which they are incapable of doing. However, on both the average ability and heterogeneous accounts, we are demanding of people with below average cognitive capacities that they perform beyond their capacities. If we take "ought" implies "can" seriously, then neither of the latter two accounts of epistemic duty are coherent. However, if we give up "ought" implies "can", it is unclear what reason we have, apart from our intuitions, about cases to support either of the two latter accounts.

If we take fitting-ness with our intuitions about cases to be sufficient reason to select one conception of epistemic obligation over another, then the heterogeneous account seems better since it fits most of our intuitions about cases. However, this account seems too ad-hoc. There is no adequate explanation as to why those of below average cognitive capacities ought to reason more and better than they are capable of but those of average capacities have no such obligation. This suggests that there is no coherent principle that could explain all three intuitions. If so, we have to give at least one of them up.

It is absurd to think that the most rational among us are those who are least capable of reasoning. It is equally absurd that the most capable reasoners among us are not justified in believing the conclusions that they have reasoned impeccably towards and which are supported by their evidence. It is, however, not so absurd to think that something has gone wrong when we suspend judgment about a proposition even when our evidence supports it just because the connection between evidence and proposition is too complicated. To admit, at least, that our cognitive capacities or resources are limited is to concede that we are imperfect reasoners. Insofar as we are imperfect, we cannot, when those imperfections affect our cognition, regard our beliefs as fully justified. Distinguishing praiseworthy or blameworthy cognition from justified or rational cognition goes some way towards ameliorating the incongruity of saying that people are not justified in suspending judgment when the evidence becomes too complicated. Of all the intuitions which we should give up on, it seems that we should concede that it is not fully rational to suspend judgment in cases where our evidence supports a proposition, but our cognitive capacities or resources are too limited to determine that it does.

If the above considerations are right, there is no plausible Forgiving account of rationality. Therefore, we should accept Simple and hence the Naïve account of epistemic duties. Recall that on Simpson's account, Nancy Drew may permissibly ignore the sincerity of testimony because assessing it is not within her cognitive capacities. Similarly, Veronica Mars may permissibly assess her evidence piece by piece because evaluating it as a whole is beyond her capabilities. Their disagreement, on Simpson's account, counts as rational because they may permissibly fail to take account of certain considerations within their purview. If the Naïve account is true, then their disagreement does not count as rational and is therefore not an instance of permissive disagreement.

To sum up, Simpson's account of how different agents can have different epistemic standards relies on each agent finding it difficult to respond optimally to all of their reasons. It is for that reason that their response does not count as rational given our commitment to full rationality. Naturally, actual moral agents are limited, however, our concern is not with what would be the wisest course of action when deciding how to assess their evidence, but with what reasons they have and what their reasons require. Assessing what their reasons require means assessing what all of their evidence implies with regards to a given proposition. Epistemic standards which are indexed to the agent's cognitive capacities would not deliver rational responses to the evidence, unless the agent can reliably respond well to all aspects of the evidence. However, there is no reason to think, on Simpson's picture, that flawless agents will have different epistemic standards. Therefore, Simpson's account of permissivism does not succeed.

Taking a step back, the claim that the first, third and fifth burdens of judgment were sources of rational disagreement ultimately relied on the claim that people could permissibly lower their epistemic standards in response to their limited cognitive capacities. In showing that they could not permissibly do so, I have shown that the first, third and fifth burdens of judgment could not possibly be sources of rational disagreement. In doing so, I have also closed off one possible avenue by which people can have different epistemic standards and hence shown that one potential way in which the disagreement generated by the second burden rational. All that is left are two ways in which people having different backgrounds could be a source of rational disagreement. In the next chapter, I shall discuss a version of permissivism according to which the epistemic standard that is appropriate for an agent depends on the relative extents she values acquiring truths and avoiding falsehoods.

Chapter 6: Indexing Epistemic Standards to Agent's Valuations of Their Epistemic Goals

Background

In the previous chapter, I discussed the possibility that the epistemic standard that was appropriate for a given agent depended on her cognitive capacities. I argued that such an indexing proposal was not successful because it presupposed that people could permissibly lower their standards just because they were unable to reliably respond to their evidence. An alternative indexing proposal is that which epistemic standard is appropriate for a given agent depends on how she might value any given epistemic or pragmatic goals. This view, which I shall call epistemic instrumentalism, treats epistemic rationality as just another species of instrumental rationality⁴³. Whether it is rational to believe a proposition given the evidence is a matter of how much the agent desires believing the truth about the proposition, and how much she desires avoiding falsity; or for that matter, how much she desires achieving any other goal by believing or not believing the proposition in question. To illustrate, consider the following example:

Religion: Sally and Susan have heard all the arguments for and against the existence of God. Sally has been raised to value avoiding false beliefs more than she values acquiring true beliefs and suspends judgment about whether God exists. Susan has been raised to value acquiring truths more than avoiding falsehoods and therefore believes that God exists.

Since Sally values avoiding falsehoods much more than acquiring truths, she suspends judgment. By contrast, Susan values acquiring true beliefs more than Sally does, and hence forms a belief that God exists. As per my discussion of the fourth burden in Chapter 2, one way in which people's different backgrounds is supposed to generate rational disagreement is by them valuing acquiring truths and avoiding falsehoods differently. In Religion, it is these different valuations which

⁴³ William James (1896) offers an account of instrumentalist permissivism in his argument for the religious hypothesis. Kelly (2014), Kvanvig (2014) and Levinstein (2017) seem to suggest that at least some version of this Jamesian permissivism might be plausible even if they stop short of a full endorsement. See also Pettigrew (2016) for a formalisation of this Jamesian approach, and Horowitz (2018) for an argument against the Jamesian approach. Horowitz's argument is that only fully proper scoring rules are consistent with updating beliefs by Bayesian conditionalization and all fully proper scoring rules recommend the same doxastic attitude given the same evidence and priors. Whereas Horowitz's argument relies on the assumption that we ought to update credences by conditionalization on our priors, the argument I offer here makes no such assumption.

supposedly give rise to rational disagreement between Sally and Susan. Notably, what is presupposed in such an account of permissivism is that these different ways of valuing these goals are rationally permissible. If only one way of valuing truths and falsehoods is permissible or of the different permissible ways of valuing these goals they all recommended the same doxastic attitude, then at least one of Susan or Sally must be making a mistake and hence be irrational.

The argumentative strategy of this chapter is to show that only certain ways of valuing these goals can be permissible and these ways of valuing these goals all give rise to the same epistemic standard. If this is right, even if agents' different backgrounds made them value acquiring truths and avoiding falsehoods differently, that would not render their resulting disagreement rational. If the argument in this chapter succeeds, then this will also illuminate whether other pragmatic goals can permissibly affect how much the agent values believing the truth about a given proposition and not being wrong about it. To be clear, instrumentalism has traditionally been defended and critiqued as an account of the source of epistemic normativity⁴⁴ (Cote-Bouchard 2015: 339). I shall not be focusing on the disputes regarding the normative force of epistemic reasons. Rather, my focus shall be on whether instrumentalism could ground a version of permissivism and I shall not address whether this has any further implications for instrumentalism.

6.1 Instrumentalist Permissivism

To start with, I shall focus on the value of the twin epistemic goals. The instrumentalist case for permissivism relies on two claims. The first claim is that there are different permissible ways in which we could value two distinct epistemic goals: believing truths and avoiding falsehoods. The second claim is that each permissible way of valuing these goals licenses a different epistemic standard. I shall explain the second claim before I discuss the first claim.

Recall that for Schoenfield, "the agent has high confidence that forming opinions using her standards will result in her having high confidence in truths and low confidence in falsehoods" (Schoenfield 2014: 199). As I have discussed in the fourth chapter (4.1.4), agents have to make trade-offs between having high confidence in truths and low confidence in falsehoods. For instance, if one is comparatively "bolder" and tends to form higher confidence in propositions based on the evidence,

⁴⁴ For recent defenders of instrumentalism, see Foley (1987), Kornblith (1993, 2002), Laudan (1990a; 1990b), Maffie (1990), Nozick (1993) and Papineau (1999). For some recent criticisms see Kelly (2003), Leite (2007), Lockard (2013) and Berker (2013).

one risks not only having a very high confidence in false propositions but also having similarly high confidence in truths. On the other hand, if one is more “cautious”, one will have more centrally distributed confidences. One might have fewer very high confidences in truths, but might also have much fewer very high confidences in falsehoods. Different ways of weighting these epistemic goals will license different epistemic standards. As I have suggested in the second chapter, one way in which, according to the permissivist, people’s different backgrounds might be a source of rational disagreement is that those different backgrounds make it permissible for agents to weight these epistemic goals differently. Perhaps valuing avoiding falsehoods is a norm in the culture one was raised in, or is a cognitive disposition prevalent at one’s workplace. Only if these different weightings are appropriate for the agent would it likewise be that the resultant epistemic standards are rationally permissible. The task in this and the next sections (6.1 and 6.2) is to show that while some of these weightings of these epistemic goals are appropriate, all of them result in the same epistemic standard. The section that follows that (6.3) will extend the argument to pragmatic goals.

As Schoenfield explains in a footnote (2014: 214), truth-conduciveness is measured by a scoring rule, which in turn is just a function that assigns an epistemic utility value to the doxastic attitudes of propositions when they are true or false. An agent’s scoring rule represents the way in which she values the twin epistemic goals of acquiring truths and avoiding falsehoods. We can represent a scoring rule as a function, $v(A,S)$. “A” is a variable representing the agent’s doxastic attitude towards a given proposition while S is a variable representing the state of the world or truth value of that proposition. If our doxastic attitudes consisted of just belief, disbelief and suspension of judgment, A would take on one of three possible values B, D and Sj respectively. Likewise if our doxastic attitudes were represented by credences, A would take on some value between 0 and 1. S takes on the value T when the proposition is true and F when it is false. An agent acquires truths to a greater degree if she has a higher confidence in a true proposition. She avoids falsehoods to a greater degree by not having a low confidence in a true proposition. In forming a higher confidence in a given proposition, an agent thereby risks failing to avoid a falsehood in order to acquire a truth. By contrast, in forming a more moderate attitude, the agent avoids falsehoods while giving up the chance to acquire truths. On the instrumentalist picture, which attitude is rational will depend at least in part on the relative extents to which the agent values avoiding falsehoods and acquiring

truths. That is to say, the question of which attitude is rational depends at least in part on the scoring rule $v(A,S)$.

The function, $v(A,S)$, maps A and S onto a real number. That is to say, for each state of the world, it assigns a numerical value to a doxastic attitude. Following Greaves and Wallace (2006: 459), I shall suppose that $v(A, S)$ increases as the doxastic attitude approaches truth and moves away from falsehoods⁴⁵. The degree to which an agent values acquiring truths and avoiding falsehoods will affect the scoring rule. For instance, suppose we were to set the value of disbelieving a true proposition $v(D,T)$ to 0. Then, if an agent, Sally, assigns 10 units of value to acquiring truths and 30 units to avoiding falsehoods, then $v(S_j,T) = 30$ and $v(B,T) = 40$ for her. This is because, when she suspends judgment, she has avoided falsehood, but not acquired truths. When she believes a true proposition, she has both avoided a falsehood and acquired a truth. By contrast, suppose Susan assigns 10 units of value to avoiding falsehoods and 30 units to acquiring truths, then $v(S_j,T) = 10$ and $v(B,T) = 40$ for her. In both cases, $v(A,T)$ increases as the doxastic attitude approaches truth and moves away from falsehoods. However, the differences in the value of suspending judgment may license different doxastic attitudes under certain circumstances. To see why, let us extend the Religion case and consider how Sally and Susan might decide between believing and suspending judgment about the existence of God.

To make such a choice, on the instrumentalist picture, agents apply decision rules. For instance, one decision rule is to maximise the expected epistemic utility of a doxastic attitude $E(A)$.

$$E(A) = L \cdot v(A,T) + (1-L) \cdot v(A,F)$$

The expected epistemic utility of a doxastic attitude A is the value of A when the proposition is true, $v(A,T)$, multiplied by the likelihood, L , that it is true, added to the value of A when it is false, $v(A,F)$ multiplied by the likelihood, $1-L$, that it is false. This rule assumes that the agent weights losses as much as gains in epistemic utility. Agents with different appetites to epistemic risk may adopt different decision rules. To simplify the presentation of my argument, I shall assume that maximising expected epistemic utility is the correct decision rule. Nothing I say in my argument

⁴⁵ There are alternative conventions about scoring rules. For instance, Pettigrew (2016) uses an inaccuracy score. The further an attitude is away from the truth and the closer it is to falsehood, the larger the inaccuracy score. In such a case, desirable states of affairs have low inaccuracy scores and the goal, correspondingly is to minimise expected inaccuracy.

hangs on whether this, or some other decision rule, or a plurality of decisions rules is correct.

Returning to the Religion case, suppose that the overall evidence only slightly supported the proposition that God exists. To provide some illustration of the degree of evidential support, let us stipulate that propositions which are supported only slightly tend to be true about 60% of the time. Therefore, for this case we can set L at 0.6. For Sally and Susan, the value of believing a true proposition, $v(B,T)$, is 40 and that of believing a false one, $v(B,F)$, is 0. For Sally, the value of suspending judgment about a true proposition $v(Sj,T)$ should be the same as suspending judgment about a false one $v(Sj,F)$ and is set at 30. Likewise for Susan, the value of suspending judgment $v(Sj,T) = v(Sj,F) = 10$. We can now calculate the expected values of believing or suspending judgment about the existence of God for Sally and Susan. For Sally, the expected value of suspending judgment $E(Sj)$ is 30 while the expected value of believing that God exists $E(B)$ is 24. Therefore, for Sally, since $E(Sj) > E(B)$, she should suspend judgment. For Susan, $E(Sj) = 10$ and $E(B) = 24$. Therefore, since $E(B) > E(Sj)$ for Susan, she should believe that God exists.

We can attempt to generalise the link between the valuing of various epistemic goals and the scoring rule. For Sally, who valued avoiding falsehoods more than acquiring truths, $v(A,T)$ rose sharply at first then gradually later. By contrast, for Susan, $v(A,T)$ rose gradually initially and then more sharply later. We might generalise beyond the coarse grained account of doxastic attitudes and consider any three adjacent doxastic attitudes A_k , A_{k+1} and A_{k+2} where A_k represents a lower degree of confidence than A_{k+1} and A_{k+2} respectively.

Whenever an agent values avoiding falsehoods more than acquiring truths,

$$v(A_{k+1},T) - v(A_k,T) > v(A_{k+2},T) - v(A_{k+1},T).$$

If an agent values avoiding falsehoods just as much as acquiring truths,

$$v(A_{k+1},T) - v(A_k,T) = v(A_{k+2},T) - v(A_{k+1},T).$$

Likewise, if she values avoiding falsehoods less than acquiring truths,

$$v(A_{k+1},T) - v(A_k,T) < v(A_{k+2},T) - v(A_{k+1},T).$$

We can also see from this how different ways of valuing the twin epistemic goals ultimately generate different epistemic standards. Sally's scoring rule is determined by the values she attaches to the twin epistemic goals. That scoring rule, together

with her evidence, determines the expected value of suspending judgment or believing. On the instrumentalist picture, rationality is just a matter of maximising expected value. This generates a set of epistemic norms which dictate which attitude is justified given the evidence.

The above account of permissivism presupposes that people can permissibly value the twin epistemic goals differently. I shall argue that, for all but one of the possible epistemic standards formed on the basis of instrumental considerations, these standards commit the agents who adopt them to two inconsistent beliefs.

To see why, consider what it takes, on the instrumentalist account for an agent to pick one doxastic attitude rather than another. An epistemic agent who is an instrumentalist needs, for each attitude, A , towards a proposition, the value of that attitude when the proposition is true, $v(A,T)$ and its value when it is false, $v(A,F)$. She also needs an attitude-independent assessment of the likelihood of the proposition being true or false. The first of these requirements is satisfied with assigning some real number to each attitude when the proposition is true and another for when the proposition is false. A mere ordering of attitudes is insufficient as it is true of any plausible valuation that a higher confidence in a proposition is preferable to a lower one if the proposition is true but not otherwise. A more precise specification is needed of how much more preferable one attitude is to another.

To illustrate, consider again the Religion case where we were determining which doxastic attitude was rational for Sally. If we did not assign numerical values to $v(S_j,T)$, $v(S_j,F)$, $v(B,T)$ and $v(B,F)$, we could not determine which of $E(B)$ or $E(S_j)$ was higher. If it was indeterminate as to which attitude had a higher expected value, then it would not have been possible to determine that Sally should suspend judgment about the existence of God. However, for the instrumentalist, given a scoring rule and decision rule, it cannot be indeterminate as to which attitude is rational for an agent.

The second requirement is satisfied only by some numerical specification of probability. By a numerical specification of probability, I mean that the likelihood should be specified by a number between 0 and 1 like 0.8 or 0.9 or perhaps an interval of such numbers $[0.8, 0.9]$. Suppose that agents instead used non-numerical vague descriptors like “very likely” or “somewhat unlikely”. Those agents would not be able to determine which attitude had the highest expected epistemic value. To see why, consider again what it takes for Sally to rationally suspend judgment about the existence of God. If all we can say about the proposition that

God exists is that it is somewhat more likely than not, the question of whether, given her epistemic standards, believing or suspending judgment is rational for her is indeterminate. For instance, if the likelihood is 0.6, then the expected value of suspending judgment is higher. If the likelihood is 0.76, the expected value of believing is higher. Even if the likelihood itself was imprecise, that would still give a determinate answer as to what the rational attitude was. However, the same claim cannot be made when the likelihood is spelled out in non-numerical form. There are many numerical specifications of likelihood that might be equivalent to "somewhat more likely than not". Not all such specifications, even if imprecise, recommend the same response. If the likelihood was $[0.6, 0.7]$, suspending judgment would still have a higher expected value. On the other hand, if the likelihood were $[0.66, 0.76]$, some imprecise attitude $[S_j, B]$ might be required instead.

Non numerical specifications of likelihood are too ambiguous to pick out one doxastic attitude given an epistemic standard. In short, what is required is a numerical representation of the likelihood of some sort or another. If no numerical likelihood is provided, this might seem to license both suspending judgment and believing that God exists.. However, this violates intrapersonal uniqueness; given an epistemic standard, at most one doxastic attitude towards a proposition would be a rational response to the evidence. Therefore, both suspending judgment and believing cannot be licensed by the evidence, given Sally's epistemic standard. Intrapersonal uniqueness therefore requires that the likelihood be given a numerical value.

However, if agents are committed to attributing a numerical likelihood to a proposition, they will, in all but at most one case, be committed to attributing another likelihood to the same proposition. This would be because doxastic attitudes can be associated with probabilities as well. Take an agent, Sally, who suspends judgment about whether God exists. Sally is committed to associating some probability with suspending judgment. The reason for this is that it can always be asked of Sally how often she expects propositions, about which she suspends judgment, to be true. If the likelihood is really low, she would be surprised if the proposition turns out to be true. If the likelihood is high, she would be surprised if it turns out false. It seems generally true that we expect propositions which we have a higher confidence in to be more likely to be true than those which we have a lower confidence in. It is hard to see why it could be permissible to attribute an equal or higher chance to a proposition we had less confidence in. Similarly, more imprecise doxastic attitudes should also elicit imprecise chance attributions from us. If the

foregoing is correct, then there is some likelihood associated with suspending judgment and a different likelihood associated with believing a proposition. Suppose that Sally and Susan expect propositions they suspend judgment about to be true roughly half the time. To be somewhat more precise, let us suppose that they expect propositions which they suspend judgment about to be true between 35% to 65% of the time, and propositions that they believe to be true more than 65% of the time. Associated with her doxastic attitude, Sally is committed to attributing a likelihood between 0.35 and 0.65 to the existence of God while Susan is committed to attributing a likelihood between 0.65 and 1 to the same proposition. Moreover, both Susan and Sally are committed to attributing a likelihood 0.6 to the proposition that God exists. While these two attributions are consistent for Sally, Susan's are not. It cannot be that the likelihood of God's existence is both 0.6 and between 0.65 and 1. Since the way Susan values her epistemic goals commits her to attributing inconsistent likelihoods to the proposition that God exists, her valuation of the epistemic goals must be impermissible.

To generalise, Sally, on the basis of her epistemic standard S_1 , forms attitude A_1 about a proposition P in response to the evidence. Given instrumentalism, she is able to do this because she antecedently attributes a numerical likelihood L to P being true. Given her attitude A_1 , she is also committed to a certain likelihood L_1 of P being true. L_1 may or may not be identical to L . If L is different from L_1 , then she is committed to inconsistent beliefs about the likelihood of P . She can assign either L_1 or L to P being true, but she cannot rationally assign both L_1 and L to the truth of P . The only way to escape irrationality is for L_1 to be identical to L . From this, we can see why uniqueness has to be true.

I shall assume that different epistemic standards, S_2 to S_n , recommend different attitudes, A_2 to A_n . It follows that the associated probabilities L_2 to L_n will be different from one another and from L_1 . Therefore, among all the likelihoods L_1 to L_n , at most one of them would be identical to L . All other epistemic standards would commit Sally to some likelihood other than L in addition to L . Therefore, at most one epistemic standard would yield a rational doxastic attitude.

It is important to emphasise that on the instrumentalist picture, Sally could not assign any likelihood to P other than L . Suppose that instead of 0.6 there was some other numerical likelihood 0.8 that Sally could assign to the likelihood that God exists.. It is clear that Sally could not permissibly choose between 0.6 and 0.8. Or else, choice of 0.6 or 0.8 would be arbitrary. It is equally clear that given that Sally

permissibly accepts 0.6, there is no reason why anyone else like Susan would have any basis for choosing 0.8 over 0.6. It cannot be because they have different epistemic standards. According to the instrumentalist, an agent's epistemic standards arise from her scoring rules, i.e. how she values particular attitudes towards propositions when they are true or false, and her decision rule, i.e. how she accounts for the risk of those propositions being false. If the rationality of a doxastic attitude is to be sensitive to the evidence, then L should be generated in some way by the evidence, independently of the agent's epistemic standards. If the assignment of a likelihood L is prior to her epistemic standards, then any scenario in which Sally is able to choose between L and another likelihood L* is subject to the Arbitrariness Objection. After all, Sally could just as easily come to a different belief that was equally consistent with one's own epistemic standards by popping a pill as by adopting L* instead of L. Therefore, at most only one of L or L* is appropriate for Sally to adopt. If the choice of L or L* is independent of which epistemic standards the agent accepts, then at most one of L or L* is appropriate for everyone to adopt.

If assigning L to P being true was dependent on epistemic standards, then epistemic standards cannot be a matter of scoring rules and decision rules. Given that the correct decision rule is to maximise expected utility, there is at most one set of scoring rules which generates doxastic attitudes which commit an agent to L. Other scoring rules will commit her to different likelihoods of P being true as well as L. More generally, the only acceptable combinations of scoring and decision rules will be those that generate doxastic attitudes that commit an agent to L. Since we differentiate between epistemic standards only by whether they generate different doxastic attitudes, these different combinations of scoring and decision rules still amount to the same epistemic standard. Therefore, if epistemic standards are to determine which of L, L* or any alternative is to be assigned to P being true, epistemic standards cannot be determined by scoring rules and decision rules. That is to say, epistemic standards cannot be determined by the extent to which we value acquiring truths and avoiding falsehoods or how cautious or daring we are in risking being wrong about a proposition.

6.2 Deciding Without Presupposing a Likelihood

The above argument relied on the premise that in order for an agent to choose between doxastic attitudes towards a proposition P, she was committed to attributing some numerical likelihood L to whether P. The reason for this rested on the claim that any possible decision rule has to presuppose some likelihood of P.

However, strictly speaking, this is not true. At least some decision rules under conditions of uncertainty (where probabilities are unavailable) do not presuppose any given likelihood that P . The most obvious examples are the maximin and maximax decision rules.

To explain these decision rules, I shall return to the earlier discussed Religion case whereby Sally is choosing between suspending judgment and believing that God exists. In this example, given her epistemic standard, suspending judgment had a value of 30 regardless of whether it was true or false and believing that God exists has a value of 40 when true and 0 otherwise. On the maximin rule, Sally should choose the attitude which maximises the prospects of the worst outcome. Therefore, on this rule, she suspend judgment since the worst outcome for suspending judgment, 30 units of value, exceeds the worst outcome for believing, 0. On the maximax criterion, Sally should choose the attitude in which the best outcome has the highest value. This means that she should believe that God exists instead since the best outcome for belief is worth 40 units of value while the best outcome for Suspending judgment is 40. The maximax and maximin decision rules do not presuppose any given likelihood of the proposition that God exists since they deliver the same verdict regardless of which likelihood we may assume to be the case. For instance, if the evidence suggests that the likelihood of God's existence is 0.99 instead of 0.6, the maximin rule would still require Sally to suspend judgment while the maximax rule would still require her to believe.

It should be clear enough that if we fix the decision rule, no plausible scoring rule would yield a different decision. If God exists, then the attitude believing is closer to the truth than suspending judgment. That is what explains why the value of believing is higher than the value of suspending judgment when the proposition in question is true. Any plausible scoring rule will value believing more than suspending judgment about a true proposition. The maximax decision rule will therefore always pick belief over suspension of judgment. Since belief is closer than suspending judgment to the truth if God exists, the reverse must be the case if He did not. Therefore, on any plausible scoring rule, suspending judgment would be better than believing given that the proposition in question is false. The maximin decision rule will always pick suspending judgment over believing regardless of which scoring rule the agent accepts. If fixing the decision rule is sufficient to fix which attitude the agent ought to pick, then given either of these two decision rules, the epistemic standard of an agent cannot depend on the agent's scoring rule. That

is to say, if there is only one correct decision rule, then only one epistemic standard is rationally permissible.

The permissivist must therefore claim that agents can have different epistemic standards because they have different decision rules. The maximin rule represents a maximally risk averse disposition. This is because when applying the maximin rule, the agent pays attention only to the worst outcome. By contrast, the maximax rule represents a maximally risk loving disposition. Only someone who is not concerned about any outcome other than the best outcome would apply the maximax rule. Like this, we are able to set up a permissive case: One agent, Sally is maximally risk averse and uses the maximin rule. Therefore she employs epistemic standard S1 and suspends judgment in response to her evidence. Susan, another agent, is maximally risk loving and adopts the maximax decision rule. Therefore she employs standard S2 and believes that God exists in response to her evidence.

The problem with this picture of permissivism is that it faces a dilemma. Either agents are permitted to be implausibly insensitive to the evidence or there are possible situations in which the range of permissible doxastic attitudes is implausibly gappy.

The first horn of the dilemma is especially obvious when the only two doxastic attitudes available to the agent are suspending judgment and believing a proposition. As mentioned in the case where Sally suspends judgment about the existence of God because she adopts the maximin decision criterion, she would suspend judgment about the existence of God regardless of how likely it was given her evidence that God exists. Moreover, this would be true for all the propositions that she would consider. Suspending judgment is always the least risky option. But this seems implausible. It cannot be permissible for Sally to suspend judgment about all propositions regardless of how well supported they are. Likewise, when Susan applies the maximax decision rule, she will believe or disbelieve all propositions that she considers and will never suspend judgment. However, it is implausible that it would be permissible for an agent to never suspend judgment. Surely there are at least some cases where the evidence for and against a proposition are so exactly balanced that suspending judgment becomes the only rational response.

When the available doxastic options are so coarse grained, adopting doxastic attitudes on the basis of the maximin and maximax rules seems unduly insensitive to the evidence. However, this changes when more doxastic attitudes become

potentially available⁴⁶. Suppose that instead of just believing or suspending judgment, agents could suspend judgment, moderately believe or strongly believe a proposition. Then, if the evidence was somewhat stronger, perhaps only moderately and strongly believing a proposition might be available to the agent. Sally, who applied the maximin rule, would moderately believe the proposition while Susan, who applied the maximax rule, would strongly believe the proposition. By contrast, if the evidence was somewhat weaker only suspending judgment and moderately believing the proposition would be available to the agent. In this case, Sally, who applied the maximin principle, would suspend judgment while Susan, who applies the maximax principle would moderately believe the proposition. In this case, unlike when only suspending judgment and believing are available, Susan's and Sally's doxastic attitudes exhibit some significant degree of responsiveness to the evidence.

However, this opens up the possibility that the range of doxastic attitudes that are permissible for someone or another, given the evidence, is gappy. To see why, consider the following version of the Religion case.

Religion II: Susan and Sally are considering whether God exists. Their evidence favours the proposition to some degree. Available to them are three doxastic attitudes, suspending judgment, moderately believing and strongly believing. Sally was raised by sceptics who are highly averse to epistemic risk and applies the maximin rule. Therefore, she suspends judgment. Susan was raised by Jamesian pragmatists who have a large appetite for epistemic risk and applies the maximax rule. As a result, she strongly believes that God exists.

If the maximax and maximin decision rules were the only rationally permissible decision rules, then of the available doxastic attitudes, there would be a gap between the permissible doxastic attitudes. Strongly believing and suspending judgment about God's existence would be permissible, but moderately believing that God exists would not.

However, it is implausible that there could be such gaps in the set of permissible doxastic attitudes. A key permissivist intuition is that it is implausible, especially with regards to inductive and abductive inferences, that it can be rational to have a

⁴⁶ Presumably, not every possible doxastic attitude an agent may adopt is available to her. For instance, if the evidence favours a proposition to some degree or other, disbelief is closed off to her.

certain degree of confidence in a proposition given a body of evidence, but not rational to have a slightly lower or higher confidence. Kelly (2014: 299-300) argues that intuitively, especially with respect to more fine grained doxastic attitudes, it cannot be that just one and only one attitude is rational given the evidence. Rather, it must be the case that there is a range of permissible attitudes which are mutually exclusive and adjacent. Moreover, whereas permissivists (ibid) are willing to concede that at least some cases are not permissive, it is less clear that they may accept that there can be cases in which the range of permissible doxastic attitudes is gappy.

In order to resist the claim that there are gaps in the range of permissible doxastic attitudes, it needs to be the case there are permissible decision rules which reflect a risk appetite between maximal risk aversion and being maximally risk loving. Moreover, it needs to be the case that at least some or other of these intermediate decision rules will recommend doxastic attitudes that will fill in these gaps. However, such a decision rule may not be available in every case. This is ultimately because any intermediate decision rule must presuppose some likelihood, L , that the proposition in question is true.

To see why this could be the case, consider the Hurwicz criterion which Pettigrew (2016) employs in his formalisation of James's claims. For our purposes, we might define the Hurwicz criterion as a decision rule which requires us to maximise the value of a function $H(A)$ where $H(A)$ is the weighted sum of the $v(A,T)$ and $v(A,F)$.

If $v(A,T) > v(A,F)$:

$$H(A) = (1-\lambda) \cdot v(A,T) + \lambda \cdot v(A,F)$$

While if $v(A,T) \leq v(A,F)$:

$$H(A) = \lambda \cdot v(A,T) + (1-\lambda) \cdot v(A,F)$$

The λ symbol is the coefficient of pessimism. It weights the value of the worst outcome and represents the degree to which the agent is risk averse. When $\lambda=1$, the agent is maximally risk averse while when $\lambda=0$, she is maximally risk loving. Consider the Hurwicz criterion: When $\lambda=1$, the criterion reduces to the maximin decision rule, whereas when $\lambda=0$, it reduces to the maximax rule. Considering the way in which the Hurwicz criterion seems to be a generalised decision rule for decision-making under conditions of uncertainty, it might seem that the Hurwicz criterion does not presuppose any particular likelihood of P . However, this would be

mistaken. The Hurwicz criterion implicitly claims that P is just as likely as not-P. To see why this is the case, consider first the explicitly risk neutral expected utility criterion.

$$E(A) = L \cdot v(A, T) + (1-L) \cdot v(A, F)$$

The agent's weighting of the two utilities by their corresponding likelihoods cannot be taken to be a departure from risk neutrality. Abstracting from the value of L, the agent who employs the "maximise expected utility decision rule" cares just as much about the worst outcome as the best outcome. That is to say, implicit in the expected utility calculation is a λ -value of 0.5. Similarly, abstracting from the weighting of the utilities in the Hurwicz criterion by λ , there must be an implicit L-value of 0.5 as well.

Secondly, consider that since each of L and λ is an independent consideration we might decide to weight the relevant epistemic utilities by, we could generate a criterion where the utilities are weighted by both L and λ . Consider, for instance, the Risk Weighted Expected Utility criterion (RWE):

$$RWE(A) = (1-\lambda) \cdot L \cdot v(A, T) + \lambda \cdot (1-L) \cdot v(A, F)$$

Consider the decision rule that maximises RWE. If we set λ to 0.5, it becomes equivalent to the risk neutral maximising expected utility rule. When we set λ to 0, it reduces to the maximax rule and to the maximin rule when λ is set to 1. This is a consequence of the fact that when the other term disappears, the value of L (or 1-L) becomes irrelevant to the rule. Crucially, when we set L to 0.5, the maximise RWE rule becomes equivalent to the Hurwicz criterion. This suggests, strongly, that the Hurwicz criterion presupposes that likelihood of P is 0.5. If this is right, then any decision rule that represents some intermediate appetite towards risk presupposes some value of L.

It follows that the only permitted intermediate decision rules are those, which in combination with the agent's scoring rule would recommend a doxastic attitude which is associated with a likelihood consistent with L. To illustrate this point, let us return to the Religion II example. In Religion II, Sally applied the maximin rule and suspended judgment about the existence of God. Susan applied the maximax rule and had a strong belief that God exists. This left a gap wherein a strong belief and suspending judgment were permissible, but moderate belief was not.

If that gap can be filled, it is possible that there is some other agent, Sharon, who adopts the maximise RWE rule at some given intermediate value of λ . For illustrative purposes, let us suppose that $L = 0.6$ in this case. In order to determine which attitude is rational for Sharon, let us stipulate Sharon's scoring rule and λ value in the following way. Let the value of suspending judgment, both when the proposition is true, $v(Sj,T)$, and when it is false $v(Sj,F)$ be 30. The value of a strong belief in a true proposition $v(Sb,T)$ is 40 while the value of that same attitude when the proposition is false $v(Sb,F)$ is 0. The value of a moderate belief in a true proposition $v(Mb,T)$ is 36 while the value of that moderate belief when the proposition is false is 20. Let the value of λ be set to 0.4. Calculating Sharon's RWE for each available doxastic attitude, $RWE(Sj) = 15.6$, $RWE(Mb) = 16.16$ and $RWE(Sb) = 14.4$. Therefore, a moderate belief in the existence of God maximises RWE.

However, it does not follow that a moderate belief is permissible for Sharon. Suppose that Sharon expects propositions which she has a moderate belief in to be true somewhere between 65% to 85% of the time. This commits her to attributing a likelihood to the existence of God that is between 0.65 and 0.85. However, she is also committed to attributing a likelihood of 0.6 to the proposition that God exists. These two attributions are not consistent with each other. Therefore, a moderate belief cannot be rational.

If, on the other hand Sharon was slightly more risk averse and had a λ value of 0.5, her $RWE(Sj)$ would be 15, $RWE(Mb)$ would be 14.8 and her $RWE(Sb)$ would be 12. Therefore, suspending judgment would maximise RWE. Given that she expects propositions she suspends judgment about to be true somewhere between 35% to 65% of the time, she is committed to attributing a likelihood to the existence of God that is between 0.35 and 0.65. Sharon's doxastic attitude therefore counts as rational. However, there is still a gap. Sharon and Sally permissibly suspend judgment while Susan permissibly has a strong belief.

It seems that no one can permissibly have a moderate belief in the existence of God. The reason for this is that any decision rule which would recommend a moderate belief must presuppose some value of L , which in this case is 0.6. However, this is inconsistent with the range of probabilities associated with having a moderate belief.

To be clear, if the evidence was such that L was 0.7, then there would be no gap. Any intermediate decision rule which recommended a moderate belief would be

rational since there would be no inconsistency between the two likelihood attributions. However, the point here is not that it is impossible to fill a gap; only that there must be some situations where there is a gap in the set of permissible doxastic attitudes.

Summing up, with regards to the maximin and maximax decision rules, when the only available doxastic attitudes are suspending judgment or belief, forming attitudes on the basis of these decision rules is implausibly insensitive to the evidence. On the other hand, when the available set of doxastic attitudes expands, the set of permissible attitudes is gappy. Given that both results are implausible, the only way to avoid either horn of the dilemma is to regard the maximin and maximax decision rules as impermissible.

This reflects the thought that the possibility of anyone adopting the epistemic standards corresponding to maximal risk aversion and risk lovingness is merely theoretical. It is doubtful that anyone is so risk averse in their epistemic outlook that they disregard completely the possibility that they may get things right. Similarly, it is doubtful that anyone is so comfortable with epistemic risk that they disregard completely the possibility that they may be wrong. Attitudes towards epistemic risk that are this extreme typically fall into the category of epistemic vices. Cassam (2016) in his discussion of epistemic vice mentions gullibility as an epistemic vice. Disregarding completely, the possibility that one might be wrong seems to be akin to some version of gullibility, rashness or arrogance. Similarly, the extreme degree of risk aversion represented by the maximin rule plausibly entails extreme cowardice on the part of the agent. If this is right, then since such extreme attitudes towards epistemic risk are characteristic of epistemic vices, the corresponding decision rules cannot be permissible.

Without the maximin and maximax decision rules, all that remains are more moderate decision rules. As I have discussed in this section, applying any such decision rule presupposes some likelihood L . Insofar as a decision rule, in combination with the agent's scoring rule recommends doxastic attitudes that are associated with likelihoods inconsistent with L , that combination of decision and scoring rule is impermissible. Since there is only one doxastic attitude which is associated with likelihoods consistent with L , all permissible combinations of scoring and decision rules will, given the same total body of evidence recommend the same doxastic attitude. That is to say, all permissible combinations of scoring and decision rules generate the same epistemic norm. Given that epistemic standards

are to be distinguished from each other only insofar as they are constituted by different epistemic norms, all the permissible combinations of scoring and decision rules will generate the same epistemic standards. If so, then even if different backgrounds cause people to have different scoring rules and decision rules, either the combination of scoring and decision rules generates the same epistemic standard, or said combination is not rationally permissible.

Summing up so far, on the indexing proposal discussed in this chapter, the epistemic standard that is appropriate for an agent depends on the extent to which she values the twin goals of acquiring true beliefs and avoiding false beliefs. I have argued that for all except at most one epistemic standard, agents would be committed to inconsistent attributions of likelihood of the proposition under assessment.

6.3 Pragmatic Goals

The above argument also allows us to examine why different pragmatic goals would not make rational disagreement possible. Pragmatic considerations can work in a number of ways. For instance, they may affect the stakes involved in having a true or false belief about a proposition. Consider the following case:

Foetal Personhood: Sally and Susan are considering whether a foetus at 16 weeks of gestation is a person. Sally, who was raised in a feminist household, is especially sensitive to the moral risk of unduly restricting women's autonomy. Given that classifying a foetus as a person will make it inappropriate for a woman at that stage of her pregnancy to have an abortion, Sally believes that a 16 week old foetus is not a person. Susan, who was raised in a religious household is especially sensitive to the moral risk of killing a person. Given that classifying a foetus as a non-person will make it appropriate to abort it, Susan believes, given the same total body of evidence, that a 16 week old foetus is a person

In addition to the way in which the agent values acquiring truths and avoiding falsehoods, an agent's background could also affect the way in which they weigh moral and prudential risk. As we saw in Foetal Personhood, different ways of weighing the moral risks results in valuing truly or falsely believing that a foetus is a person differently.

Arguably Susan and Sally in Foetal Personhood are not so much reasoning, as engaging in wishful thinking. This seems to be an inevitable risk when the moral stakes in question are closely tied to the proposition under deliberation. Insofar as there is a dearth of plausible cases where different weighing of the moral stakes licenses different doxastic attitudes, this emphasises how implausible it is that people's doxastic attitudes towards moral claims might permissibly be influenced by the moral stakes at hand. Be that as it may, for the sake of argument, we might proceed as if Foetal Personhood or some other case presented an intuitively plausible instance of moral stakes affecting what it is rational to believe.

To see how practical considerations affect the rationality of disagreement, consider a simpler case, namely the following version of the Bank case discussed in Chapter 4.

Bank: Sally needs to urgently deposit a cheque into the bank. If she fails to deposit the cheque by the end of the day, she will lose her house. Her memory, which she knows is fairly, but not perfectly, reliable tells her that the bank will close at 7:30 pm. However, since she cannot afford to be wrong about when the bank closes, she suspends judgment about whether the bank will close at 7:30 pm and checks the internet again. Susan does not need to urgently deposit a cheque into the bank. She will lose nothing of significance if cannot go to the bank by the end of the day. Her memory which she knows is fairly, but not perfectly reliable (in fact just as reliable as Sally's) tells her that the bank will close at 7:30 pm. Since she can afford to be wrong about when the bank closes, she believes that it will close at 7:30 pm.

In Bank, Sally and Susan do not in general have different attitudes towards epistemic risk nor do they in general value acquiring truths and avoiding falsehoods differently. However, Sally's practical circumstances make it such that she happens to value not falsely believing that the bank will close at 7:30 pm much more than she values truly believing that it will. As a result, her expected epistemic value if she suspended judgment is higher than her expected epistemic value if she believed that the bank would close at 7:30 pm. That is what should make suspending judgment rational for Sally. Likewise, the opposite is true for Susan. Notice, however, that in order for it to be true that the expected value of Sally believing that the bank would close at 7:30 pm is lower than the expected value of her suspending

judgment, there needs to be some likelihood L that is assigned to the proposition that the bank will close at 7:30 pm. Without that likelihood, the expected values of believing and suspending judgment would be indefinite. There are also likelihoods associated with suspending judgment and believing. Presumably Sally and Susan expect propositions that they suspend judgment about to be true roughly as often as not. Likewise, they must expect propositions which they believe to be true significantly more often than not. However, if the likelihood of the bank closing at 7:30 pm is L , the proposition that the bank will close at 7:30 pm cannot be both significantly more likely to be true than not and roughly just as likely to be true as not. Therefore, either Sally, Susan or both must be attributing inconsistent likelihoods to the proposition that the bank will close at 7:30 pm. It follows that any epistemic standard which requires them to attribute such inconsistent likelihoods is not rationally permissible.

Having dealt with the simpler case, Bank, we can return to somewhat more complicated cases like Foetal Personhood. If it is not rational to suspend judgment about the time the bank closes just because of the high stakes of getting the timing wrong, it similarly cannot be rational to believe that the foetus is a person just because of the high moral stakes of falsely believing that it is not. Likewise it cannot be rational to believe that the foetus is not a person just because of the high moral stakes of falsely believing that it is.

To generalise, for any putative pragmatic goal, the pragmatic goal could possibly affect which attitude it would be rational to have about a proposition only by changing how valuable it is to acquire that belief if true or avoid it if false. For the same reasons⁴⁷ as with our discussion of the twin epistemic goals, in all but at most a few of the ways of valuing these pragmatic goals, this results in the agent being committed to attributing inconsistent likelihoods to the same proposition. This cannot be rational. Therefore, any disagreement that arises from valuing various goals differently cannot be a rational one.

If this is correct, then this closes off one of the ways in which, according to the fourth burden, difference in people's backgrounds are sources of rational disagreement. More specifically, even though different backgrounds give rise to different ways of

⁴⁷ I shall ignore those pragmatic goals the value of which is not sensitive to whether the proposition in question is true. For instance, if an evil demon threatened to kill everyone unless an agent believed that grass was purple, preventing everyone from dying may be some sort of pragmatic goal which does not require the agent to presuppose some L . However, that would clearly be the wrong kind of reason to believe that grass is purple. Whatever other virtue such a belief may possess, it is not epistemically rational.

valuing various practical and epistemic goals, this does not licence different doxastic attitudes towards a proposition when the evidence is shared. In addition, in showing that epistemic standards cannot vary on the basis of how agents value the twin epistemic goals or any other pragmatic goal, I have also closed off one more potential route by which the second burden of judgment could generate rational disagreement. In this thesis, I have, thus far, shown that the first, third and fifth burdens cannot be sources of rational disagreement. I have also shown that, with regards to the fourth burden, rational disagreement cannot arise because people's different backgrounds cause them to value acquiring truths and avoiding falsehoods differently. All that is left is to show that people cannot rationally disagree because their different backgrounds cause them to have different prior doxastic attitudes. In the next chapter, I shall discuss this indexing proposal according to which the epistemic standard that is appropriate for an agent depends on her prior doxastic attitudes.

Chapter 7: Epistemic Conservatism and Permissivism

I have so far discussed two accounts according to which epistemic standards might vary from person to person. The first is because people differ in terms of their cognitive capacities. The second is because people differ with regards to how much they value acquiring truths and avoiding falsehoods. I have argued that neither account is successful. I shall now discuss a third account according to which people can have different epistemic standards because they have different doxastic attitudes prior to encountering the evidence. To illustrate, consider the following case:

Abortion: Susan comes from a conservative religious background and happens to believe that abortion is impermissible. Sally comes from a secular background and happens to believe that abortion is permissible. Perhaps this is because they each have at most encountered evidence and arguments supporting their own side. When they both enter the university, they encounter all the arguments and evidence that pertain to the permissibility of abortion. It turns out, however, that the balance of reasons is insufficient for either Sally or Susan to change their views. However, Sally's belief about the permissibility of abortion adds just enough justification that she remains justified in believing that abortion is permissible. Likewise, Susan's belief that abortion is impermissible adds just enough justification to her balance of reasons that she remains justified in believing that it is impermissible.

In Abortion, after going to the university, both Sally and Susan have the same total body of evidence. However, since they initially start off having different beliefs, they are permitted to maintain their beliefs because the mere fact that they happen to believe a proposition confers some degree of justification on that proposition. This last claim is called epistemic conservatism. Some version of conservatism, together with the agent's initial set of beliefs, gives rise to epistemic norms that prescribe for each agent which doxastic attitudes are permitted for them given their evidence. That is to say, an agent's prior beliefs, in combination with more general principles about belief revision generate epistemic standards.

My aim in this chapter is to address forms of permissivism grounded in epistemic conservatism. In showing that no plausible account of permissivism can be grounded in epistemic conservatism, I shall show that even if people's different

backgrounds cause them to start off with different beliefs, as long as their total body of evidence is the same, these different starting points do not permit them to disagree. To start off with, let me first explain what I mean by epistemic conservatism.

7.1 Epistemic Conservatism

Epistemic conservatism is the claim, roughly, that an agent acquires some justification for a proposition simply by virtue of her believing it. The amount of justification provided by merely believing it need not be very much. Nevertheless, as with Abortion, we can construct permissive cases⁴⁸ whereby two agents Sally and Susan share the same evidence, where the strength of the evidence is just short of justifying belief in a given proposition.

As seen from the Abortion case it seems that so long as there is a justificatory asymmetry between maintaining a belief and adopting a belief, a permissive case can seemingly be constructed. Conservatism can generate permissive cases in other ways as well. Consider the following variation on Abortion.

Abortion II: Suppose, as per the fourth burden of judgment, two agents Susan and Sally have different backgrounds. Susan was raised in a religious conservative environment and Sally in a more secular environment. Due to their different backgrounds Sally initially has a high confidence in the proposition that abortion is permissible while Susan has a low confidence in the same. Later they both enter the university where they encounter all the arguments and evidence pertaining to the permissibility of abortion. Their initial doxastic attitudes nevertheless have some “pull” on their final doxastic attitudes. After encountering all the evidence, Sally’s confidence in the proposition that abortion is permissible is still higher than Susan’s.

The most prominent view according to which the doxastic attitudes an agent has before she encounters the evidence influences the doxastic attitude she has

⁴⁸ Gaus appeals to a version of epistemic conservatism when he defends the possibility of reasonable disagreement. He appeals to the claim that we need more reason to adopt a belief than to retain it (2011: 240). While this is not strictly speaking a matter of belief conferring justification, it is still a version of conservatism that is vulnerable to the same objections as the standard account.

subsequent to it is Bayesianism. I shall briefly explain this view before explaining how a certain permissive version of Bayesianism is related to conservatism.

For the following discussion, I shall assume that doxastic attitudes can be adequately represented by numerical probabilities. Regarding a proposition A, I shall use $P_i(A)$ to represent the agent's initial doxastic attitude towards A before encountering the body of evidence in question and $P_f(A)$ to denote the doxastic attitude towards A after encountering the evidence in question. I shall use $P(A|B)$ to denote the conditional probability of A given B. Following Kolmogorov (1956), I shall define⁴⁹ conditional probabilities in the following way:

$$\text{Conditional Probability: } P(A|B) = P(A \text{ and } B) / P(B)$$

Bayesians can be thought of as endorsing the idea that when they encounter a body of evidence E, the attitude towards A, $P_f(A)$, that they ought to form is to be related to their initial assessment of the conditional likelihood of A given E, $P_i(A|E)$, by the following principle called Conditionalization:

$$\text{Conditionalization: } P_f(A) = P_i(A|E)$$

Recall that epistemic norms either are, or imply rules of the form “form doxastic attitude A towards proposition P in response to evidence E”. Combining Conditionalization with a sentence like “ $P_i(A|E1) = 0.8$ ” gives us an epistemic norm which tells us that we should form a confidence of 0.8 towards proposition A when we encounter a body of evidence E1. Given Conditional Probability, $P_i(A|E)$ can be re-written as:

$$P_i(A|E) = P_i(E|A) \cdot P_i(A) / P_i(E)$$

This makes explicit how an agent's epistemic standards depend on her initial doxastic attitudes and her initial assessment of how likely the evidence would be if the proposition had been true. These initial attitudes and probability assessments are an agent's priors. Given Conditionalization, we can see that since $P_i(A|E)$ is proportional to $P_i(A)$, $P_f(A)$ is proportional to $P_i(A)$ as well. This can explain how, in Abortion II, Susan's and Sally's initial credences about the permissibility of abortion have a “pull” on their final credences after they encounter all the evidence.

⁴⁹ Notably this is not the only way to define conditional probabilities. For instance Cox uses conditional probabilities as a primitive and defines unconditional probabilities in terms of conditional ones.

Trivially, when people have different priors, their epistemic standards⁵⁰ will be different and the doxastic attitude that is rationalised by their evidence would be different. Different priors result in different values for $P_i(A|E)$. Combined with Conditionalization, these different priors will generate different epistemic norms. Given that epistemic standards are constituted by epistemic norms, when the epistemic norms differ, the epistemic standards are different.

The permissivist's claim then is that there are at least some cases in which agents can have different priors. There is one sense in which everyone thinks that agents can have different priors. If, prior to encountering E, they had previously encountered different evidence, they can permissibly have different priors. However, in cases like these, people have different total evidence. Uniqueness is compatible with people having different total evidence rationally disagreeing. The disagreements that matter for our purposes are disagreements when people have the same total evidence. Therefore, the cases that are of interest to us are those in which people's priors permissibly differ without there being differences in their total body of evidence. With reference to the Abortion II case, what concerns us is whether Sally and Susan can have different priors while having the same total evidence. To answer the purely causal question, people having different priors without having had different total evidence might be, as suggested by the fourth burden of judgment, a result of them having different backgrounds. The epistemic question that concerns us is whether these different priors could be rationally permissible. Therefore, we must look at whether people can differ with respect to what I will call ur-priors.

Ur-priors are those priors an agent can be said to have had before encountering any evidence at all. To be clear, I am not claiming that people are ever actually in a position where they have to pick one set of ur-priors over another. Children, on average, are presumed to be capable of forming beliefs between 1.5-2.5 years of age. However, they clearly have subjective experiences well before then. This means that actual agents can plausibly be said to encounter evidence well before they are capable of forming beliefs about various propositions. Hence, actual

⁵⁰ It is a mistake to think of the epistemic norm just consisting of Conditionalization. If the epistemic norms consisted of just Conditionalization and Conditional Probability, people with the same set of epistemic norms could have different priors, and therefore rationally disagree. Clearly, given that epistemic standards are taken to fix which doxastic attitudes are rational for an agent given the evidence, the priors themselves must be part of the epistemic norms. Since the priors themselves do not constitute any given epistemic norm, the norm must be constituted by a combination of the priors, Conditional Probability and Conditionalization.

agents' earliest doxastic attitudes are formed in some sense as a response to evidence. It seems that, in principle, these earliest responses to an agent's evidence can be evaluated.

Bayesians therefore seem committed to positing the existence of ur-priors. In principle, if rationally permissible ur-priors exist, then either the agent's doxastic attitudes could have been formed by conditionalization on those priors or they could not. If they could have been formed that way, then they can be rational. If they could not have been formed by conditionalization on those ur-priors, the attitudes cannot be rational. A corollary to this is that we can project backwards from people's current doxastic attitudes and total evidence and determine what ur-priors, if rationally permissible, could rationalise an agent's current attitudes. Our current attitudes are rational only if the ur-priors that would rationalise them are themselves rationally permissible. Moreover, Conditionalization makes it clear that the priors fix what doxastic attitude is rationalised by the evidence. It follows from this that rational disagreement given the same total evidence is possible, if and only if, there is more than one set of rationally permissible ur-priors.

Subjective Bayesians, those who believe that more than one set of ur-priors is permissible, have some difficulty accounting for how people could rationally accept ur-priors. Clearly, ur-priors are subject to rational appraisal. Even the most extremely permissive Bayesians will require that ur-priors are to be probabilistically coherent. The difficulty is in figuring out what other constraints can be imposed on the "choice" of ur-priors. Meacham (2014: 1192-1193) notes, for instance, that chance-credence principles like the Principal Principle which requires agents to form credences that match their beliefs about objective probabilities impose some constraints on which ur-priors are rationally permissible. By doing this, the range of permissible ur-priors is narrowed. However, as Meacham notes (*ibid*), chance-credence principles are not so strong that only one set of ur-priors is permissible.

What is notable about the strategy of Meacham's argument is the implicit claim that a given set of priors acquires some degree of justification as long as there are no principles that rule them out. However, this just is epistemic conservatism. Part of what supposedly makes an agent's ur-priors justified for her is that they are *hers*⁵¹. By contrast if only one set of ur-priors was rationally permissible, then, for an agent who has those ur-priors the fact that those ur-priors are hers does not contribute in

⁵¹ This has to be the case or else permissive Bayesianism is subject to the Arbitrariness Objection. An agent who did not prefer her own priors would have no epistemic objection to having her ur-priors changed for her by the press of a button if such was indeed possible.

any way to their rational permissibility. Addressing permissivism that is grounded in this justificatory asymmetry between one's own doxastic attitudes and other attitudes requires us to discuss epistemic conservatism in greater detail.

While epistemic conservatism has been subject to some recent criticism (Foley 1983; Christensen 1994; Vahid 2004), it has been endorsed by many philosophers including Quine (1951), Sklar (1975), Chisholm (1980) and more recently by McGrath (2007), McCain (2008), Poston (2012) and Podgorski (2016) in various forms. Some of these may be weaker than others. My argument in this section will be that the stronger versions of conservatism are implausible while the weaker versions do not pose any significant threat to uniqueness. In the following, I shall present what seems to be the standard version of epistemic conservatism as well as what seems attractive about epistemic conservatism.

Vahid describes the standard version of epistemic conservatism as follows.

Standard Conservatism⁵²: "Beliefs acquire positive epistemic status merely in virtue of being believed in so far as there is no good reason against them" (Vahid 2004: 107)

In clarification of the above, I shall take positive epistemic status to mean some degree of justification. On a stronger account, the epistemic status in question could be full justification. On a weaker account, the epistemic status may not even survive contact with evidence (Poston 2012).

By no good reason, I mean no good positive reason (Harman 1986: 39; Vahid 2004: 109). Conservatism would be entirely toothless if the fact that there was no positive reason for a belief counted as a good reason against that belief. As Harman and Vahid note, in order for propositions I believe to have any advantage over those that I do not, it must be permissible for me to hold onto a belief until I have positive reason to abandon it, and not just if I lack positive reason to maintain it. With these clarifications in mind, we should have some grasp of what beliefs are permitted by Standard Conservatism. I shall now proceed to explain why epistemic conservatism might be attractive.

⁵² This is basically identical to how Chisholm (1980: 551-552) describes it. Foley follows Chisholm in formulating epistemic conservatism.

7.2 Motivating Conservatism

There are three⁵³ canonical sorts of motivations for conservatism. The first sort of motivation relates to cognitive efficiency. The second relates to issues pertaining to memory retention. The third relates to evidential under-determination and sceptical worries. I shall describe each sort of motivation in turn. With regards to the first motivation, on one view, changing one's doxastic attitudes is cognitively costly and epistemic conservatism is supposed to make our doxastic decisions sensitive to these costs (Lycan 1988; Podgorski 2016). Relatedly, one can also avoid intellectual clutter by not having to keep track of the justifications for all of one's beliefs (Harman 1986). Since conservatism provides some occasions where beliefs can be retained when they otherwise could not be without having to keep track of the evidence, conservatism provides a boost to the cognitive efficiency of agents.

Crucially, conservatism is a view that is invoked in response to the cognitive limitations of actual agents. If we were not cognitively limited, we would not need to be sensitive to the cognitive costs of changing one's doxastic attitudes. Similarly, intellectual clutter would not be a problem and we would always be able to keep track of all our evidence all the time. These are similar types of considerations to those that are supposed to operate in the first, third and fifth burdens of judgment. According to the first and fifth burdens, people rationally disagree because the evidence or the subject matter is too complex. The complexity of the evidence or subject matter cannot cause disagreement unless people have limited cognitive capacities. Likewise, as per the third burden of judgment, the vagueness in the concepts we use is a result of our limited cognitive capacity to infer from our evidence whether a predicate applies in borderline cases. To the extent that, as I have argued in Chapter 5, limitations to an agent's cognitive capacities do not affect what is rational for her to believe, some of the motivation for epistemic conservatism is undermined. I shall discuss this more extensively later in this chapter.

The second sort of motivation for conservatism relates to memory retention. The core case is one in which we retain a belief even after we have forgotten the reasons for that belief. For example, suppose I were to have originally believed that all the Moai on Easter Island face inwards on the basis of reading a particular article from National Geographic magazine. I reason that given that National Geographic is generally reliable about such matters, I have good reason to believe that the Moai

⁵³ I am excluding Poston's argument (2012) for conservatism based on the perspectival nature of reasoning at least partly because Poston offers little in the way of positive argument for it.

all face inward. Suppose I were to retain this belief but forget how I arrived at this belief. There is supposed to be an intuition that it is rational for me to retain this belief even when I have forgotten my reasons for the belief. According to McGrath (2007), conservatism does a better job of explaining the intuition that it is rational to retain the belief, than competing explanations like evidentialism or preservationism for both this case and other similar cases.

Preservationism is the view according to which memory preserves the rationality of a belief (Burge 1993). Hence, according to preservationism, barring the introduction of new evidence, a retained belief is rational, if and only if the belief was rational when originally formed. However, as McGrath argues, there can be cases where the belief was originally formed irrationally but retaining the belief is rational. McGrath sketches out one such case. In his example, he supposes a case where an agent believes some fact like 'Hannibal fought in the third Punic war', on the basis of testimonial evidence she has good reason to believe is unreliable (McGrath 2007, p4). The original acquisition of the belief is irrational since the source of the testimony is someone who the agent has good reason to think is unreliable about such facts; for instance, a six year old child. Suppose that the agent were to forget how she acquired the belief. Perhaps she remembers that someone told her, but not who. McGrath argues that retaining her belief after she has forgotten her evidence is rational because the rationality of retaining her beliefs is dependent only on her current epistemic situation. Since her current epistemic situation is identical to someone who rationally acquired her belief but forgot her evidence, it is rational for her to retain her beliefs (ibid). McGrath's argument is that since preservationism cannot explain the rationality of a retained belief which was acquired irrationally, preservationism cannot be true.

The second competitor to conservatism is what McGrath calls evidentialism. By evidentialism, he means a view according to which retained beliefs are justified because even though we lack the original evidence for the belief, we possess some other evidence for the belief (ibid: 8). McGrath's objection to evidentialism is that there is no plausible version of evidentialism which could justify memory beliefs. I shall not discuss the details of McGrath's arguments here. If they are right, then since preservationism is also off the table, according to McGrath, conservatism is the only account by which memory beliefs can be justified.

The third motivation for conservatism is how conservatism handles evidence under-determination and scepticism. I shall consider evidence under-determination first.

Suppose, a scientist accepts a theory on the basis of a body of evidence, and then encounters another theory which fits the evidence just as well as the one she believes. Conservatism permits her to maintain her acceptance of that first theory. It thus solves one type of under-determination problem: Whenever someone already accepts one theory which fits the evidence, one need not reject a theory just because another one which equally fits the theory is introduced (Sklar 1975: 381-383). We can see how conservatism addresses sceptical problems as well. Our evidence is equally consistent with the proposition that we are brains in vats⁵⁴ or that we are deceived by an evil Cartesian demon and the proposition that we are not. Conservatism entails that we need not lower our confidence in the existence and character of the external world just because we become aware of alternative sceptical scenarios which equally fit the evidence.

7.3 Objections to Conservatism

Given the range of upshots of conservatism, we would expect it to be less controversial than it is.⁵⁵ However, a common criticism of conservatism is that it proves too much. Epistemic conservatism, or at least the standard version of it, is subject to a range of counterexamples where conservatism seems to license intuitively rationally defective beliefs. I shall review two types of related canonical objections to conservatism, firstly the bootstrapping objection and subsequently the extra boost objection. Foley (1983: 174-175) brings up the case in which an agent considers whether there are an even number of grains of sand on a given beach. She fails to notice that she lacks sufficient reason for this proposition and believes it anyway. Foley's objection is that conservatism seems to permit the agent to continue to maintain the belief since it now has some epistemic merit in virtue of being believed and there are no positive defeaters for it. People thus seem to be able to illicitly bootstrap their way into having rational beliefs when their evidence is lacking.

Related to the bootstrapping objection, suppose that an agent has enough evidence to justify just barely believing a proposition. Having formed the belief, conservatism seems to give the proposition some extra boost over and above the existing evidence. This extra boost licenses the agent to adopt an even higher confidence (Huemer 1999: 348; McGrath 2007: 19-20). Like the bootstrapping, this extra boost is illicit. It seems odd to think that we would get an extra boost to our justification

⁵⁴ See Harman (1973)

⁵⁵ See Foley (1983), Christensen (1994), Vahid (2004)

whenever our other reasons are strong enough to justify belief. If we were to get that extra boost, barely believing any proposition would never be justified. Every time our reasons for belief are just strong enough to justify barely believing a proposition, an extra boost would justify an even stronger belief. Moreover, the reason that this attitude would never be justified would be simply because it is the weakest doxastic attitude that still counts as belief. But surely *that* is not a good reason to think a doxastic attitude is always irrational.

7.4 Weakening Standard Conservatism

It is clear that the above two objections apply at least to Standard Conservatism. If believing a proposition confers some degree of justification to it, then agents can illicitly bootstrap themselves into having rational beliefs which were not rational for them to form. They would also gain an extra boost to any proposition which gained this degree of justification simply on the basis of being believed. If these objections applied to epistemic conservatism as a whole, they would be a devastating objection to epistemic conservatism. That is why recent attempts at defending epistemic conservatism have involved a weaker version of conservatism than Standard Conservatism (Sklar 1975; McGrath 2007; McCain 2008; Poston 2012; Podgorski 2016). I shall show that those accounts of conservatism that can successfully avoid the objections pose no threat to uniqueness. I will start from the least recent account and proceed in chronological order.

7.4.1 Differential Conservatism

In considering Sklar's account of conservatism, his objection to Standard Conservatism is on the grounds that it allows illicit bootstrapping (Sklar 1975: 376-377). Instead, he endorses the following principle which, following Vahid, I shall call Differential Conservatism (Vahid 2004: 102).

Differential Conservatism: "If you believe some proposition, on the basis of whatever positive warrant may accrue to it from the evidence, *apriori* plausibility, and so forth, it is unreasonable to cease to believe the proposition to be true merely because of the existence of, or knowledge of the existence of, alternative incompatible hypotheses whose positive warrant is no greater than that of the proposition already believed." (Sklar 1975: 378)

If Differential Conservatism is true, then there is a straightforward way to construct a permissive case. Two agents can have the same evidence, but rationally accept different propositions because each of them happens to have encountered their own proposition first. For instance, Differential Conservatism licenses different attitudes in the Abortion case. Suppose the evidence regarding the permissibility of abortion fits the proposition that abortion is permissible just as much as it fits the proposition that abortion is not. Then Susan, who encounters the proposition that abortion is not permissible first because of her background, is justified in believing so. Similarly, Sally who encounters the proposition that abortion is permissible because of *her* background is justified in believing that abortion is permissible. However, Sklar's account of conservatism is at least as implausible as Standard Conservatism because it presupposes the latter.

While Sklar denies that Differential Conservatism commits him to Standard Conservatism (*ibid*), this denial is implausible. As Vahid has argued (Vahid 2004: 104), Sklar appeals to Standard Conservatism when trying to justify Differential Conservatism. Sklar considers a permissive case one in which, one person believes a proposition while another who shares the same evidence believes an incompatible one which is equally supported by the evidence. He draws an analogy to cases where people with different evidence reasonably disagree and argues that while their evidence does not differ, their total states do (Sklar 1975: 384-385). However, this move presupposes that the difference in total state, i.e. the fact that each agent believes a different proposition, provides some justification, akin to that provided by evidence, to the respective propositions believed by each agent. However, the only difference in state between the agents is that one of them happens to believe one proposition and the other agent happens to believe another proposition. That is to say, Sklar presupposes Standard Conservatism in his defence of Differential Conservatism.

Of course, it is hard to see how Sklar could avoid appealing to Standard Conservatism. It is incoherent to claim both, that I should not stop believing a proposition just because I encounter another equally well supported but incompatible one, and that believing a proposition does not confer any justification for that proposition. After all, when asked why it is wrong to stop believing in a proposition when I encounter another evidentially equivalent but incompatible one, the only answer that would suffice is that the first proposition is more justified to me. If we say that the other proposition is just as justified as mine, it would not seem wrong to switch. However, if it is permissible to switch, then believing on the

evidence is just as good as believing by popping a pill or flipping a coin. On the other hand, if it is wrong to switch, it can only be because my proposition is more justified than the other. However, the only difference between the two is that I believe one, but not the other. Therefore nothing else, but the fact that I believe it makes up that difference in justification. However, this just is Standard Conservatism, which as we saw, is implausible.

7.4.2 Defeating Conditions

Another way in which Standard Conservatism can be weakened is by adding more defeating conditions. McGrath (2007) and McCain (2008) both try to weaken Standard Conservatism by introducing more robust defeating conditions. I shall discuss each proposal in turn. McGrath is keenly aware that Standard Conservatism is vulnerable to the bootstrapping and extra boost objections because the defeating conditions are too thin. McGrath takes Standard Conservatism to consist of something like the following: “if S believes that P, then S is rational in retaining the belief that P provided and only provided that, (i), S lacks defeaters for her belief; (ii), S lacks the materials for constructing such a defeater⁵⁶” (McGrath 2007: 20). He tries to confine conservatism to forgotten evidence cases (ibid: 21) by introducing another defeating condition: “(iii) S lacks special information about her past evidence” (ibid: 20).

However, it is unclear if the third defeating condition suffices to prevent bootstrapping cases. To see why, suppose an agent forms a belief on insufficient evidence without noticing that the evidence is insufficient. Conservatism, even with McGrath’s third condition permits the agent to maintain that belief even when the agent remembers the evidence she used to originally form the belief. The only thing that speaks against her having the belief is the fact that her evidence is insufficient. However, this is not a positive defeater, only a negative one. Similarly, there are no materials that could be used to construct a positive defeater. The question of whether McGrath’s version of conservatism permits the agent to maintain her belief in this case hangs on the third defeating condition. Since the agent does not notice that she formed her belief on insufficient evidence, she does not count as having special information about her past evidence. Therefore, McGrath’s version of conservatism permits the agent in this case to retain her belief. However, insofar as

⁵⁶ McGrath distinguishes between having a defeater and the material for the defeater being available in the following way. An agent has a defeater only if she believes that she believes the defeater. She has the material to construct the defeater, so long as she can infer the defeater from the other information she does have.

Conservatism rationalises carelessly formed beliefs, it allows illicit bootstrapping that we found to be objectionable.

McGrath could of course continue to create more defeating conditions (ibid: 21) and it is worth making explicit what it would take to ensure that conservatism applied only to forgotten evidence cases. In order to rule out everything but the forgotten evidence cases, the defeating condition would have to be something like the following: '(iv) S lacks the material to construct special information about her past evidence'. Moreover, 'special information about past evidence' would have to be interpreted broadly to include the fact that the belief lacks sufficient evidence. If special information does not include the negative claim that the evidence was insufficient for the belief, then conservatism would still permit illicit bootstrapping. Therefore the 4th defeating condition must specify, both, that the material for the special information to be constructed is unavailable, and that the special information include the fact that the belief was originally formed on the basis of inadequate evidence.

However, once we introduce the 4th defeating condition, there is pressure to regard forgotten evidence cases, in which the original evidence was inadequate, as illicit bootstrapping. To see why, consider why we would think the 4th defeating condition was appropriate. We would think that the 4th defeating condition was appropriate only because we would think it important that our beliefs be formed on the basis of sufficient evidence. If we think it important that our beliefs are formed on the basis of sufficient evidence, then it should not matter whether we currently believe that they are, only that they actually are. Even if we were to justifiably believe that our beliefs are formed on the basis of sufficient evidence, if that belief is nevertheless incorrect, then our other beliefs have at least one rational defect: They are formed on the basis of insufficient evidence. Therefore, even if we have forgotten our evidence, it is still a fact that the beliefs were formed on the basis of insufficient evidence. It would be one thing if we were to subsequently get other evidence which would then justify our beliefs. At least then even if the belief was not originally well formed, we are now in a position to justify that belief because we have more information. When our evidence was already insufficient to justify the belief, it seems absurd that taking that evidence away would render the belief justified. Therefore, if we think that the 4th defeating condition is appropriate, we have reason to think that permitting belief retention in forgotten evidence cases in which the belief was originally improperly formed also involve improper bootstrapping. However, as McGrath pointed out, taking the retention of a person's belief as irrational when she formed it irrationally

but forgot her original evidence violates the principle by which the rationality of a belief supervenes on the agent's current state.

It is true that if we do regard forgotten evidence cases in which the original evidence is improperly formed as involving bootstrapping, we are forced to give up the principle that rationality supervenes on the current state of the believer (Huemer 1999: 352; McGrath 2007: 4). However, any plausible conservatism should give up this principle anyway. Conservatism is most plausibly regarded as a principle regarding diachronic rationality, that is, about the rationality of belief revision (McGrath 2007; Podgorski 2016). Moreover, we care about diachronic rationality only because we care about more than merely whether people's beliefs fit their evidence; we care about how they formed their beliefs. If we cared only about whether agents' beliefs fit their current evidence, conservatism is dead in the water. The reason for this is that if we cared only about an agent's current evidence, we would have no reason to care about special information about the agent's original evidence. However, eliminating the third and fourth defeating conditions would make conservatism vulnerable to the bootstrapping objection. Therefore, epistemic conservatives should care about how people formed their beliefs. Insofar as we do care about how people formed their beliefs, rationality is not going to supervene on the agent's current total state. To see why, consider two agents who share the same evidence, Proper and Improper.

Proper examines the evidence carefully and forms a belief on the basis of sufficient evidence. Improper jumps to the same conclusion which just so happens to fit the evidence. Intuitively, we regard Improper's belief as being doxastically irrational while the same belief in Proper is rational. However, their current total state is the same. They have the same evidence and the same beliefs. Any difference between them is not current. Since one of them is rational but the other is not, doxastic rationality does not supervene on their current total state. To sum up, epistemic conservatism commits us to caring about how people originally formed their beliefs on pain of the bootstrapping objection. Caring about how an agent formed her beliefs commits us to rejecting the principle whereby rationality supervenes on an agent's current state. Therefore, epistemic conservatism commits us to rejecting the supervenience of an agent's rationality on her current state. If so, then, if conservatism permits retaining beliefs when they were improperly formed even though the evidence was forgotten, then, it is vulnerable to the bootstrapping objection.

It is unclear, however if, on this version of conservatism, a permissive case can be constructed. After all, the only circumstances in which the conservative principle applied would be forgotten evidence cases and even then, only when the belief was originally formed on the basis of sufficient evidence.

The reason why permissive cases cannot be constructed is that these cases all involve forgetting one's evidence and forgetting one's evidence seems like a rational defect. Even though Harman (1986), Lycan (1988) and Podgorski (2016) all acknowledge that allowing people to forget their evidence improves cognitive efficiency, this consideration is not fully epistemic. Consider that the way it improves cognitive efficiency is by allowing us to concentrate our limited cognitive resources to the task at hand or by reducing clutter. The implication here is that our limited cognitive resources, mental clutter and other sorts of cognitive limitations prevent us from responding well to all our reasons. Conservatism is a principle that permits us to fail to respond to some of our reasons so that we do a better job of responding to other reasons we have. The principle does not set a criterion for evaluating whether a given set of reasons have been fully responded to. This already concedes that in those forgotten evidence cases in which conservatism applies, agents are not fully responding to their reasons. After all, when an agent forgets her evidence but still retains the belief that was justified by the evidence, she is not responding to the evidence she has.

In this respect, conservatism, insofar as it is a strategy employed in order to cope with cognitive limitations, resembles the first, third and fifth burdens of judgment. The first and fifth burdens pertained to the difficulty of assessing the issue. While the third burden pertained more specifically to the difficulty in assessing whether a predicate applied to an object in a borderline case. This stemmed from the fact that we have limited cognitive capacities and resources. If we had infinite cognitive capacities, we would not make mistakes in reasoning and we would always be able to resolve the vagueness in our concepts. Similarly, if we had infinite cognitive resources, we would respond well to our reasons in spite of mental clutter, or our mind would have so much space it would not be cluttered. As with our earlier discussion of the first, third and fifth burdens, the fact that cognitive capacities are limited cannot justify lowering our epistemic standards.

Moreover, since an agent's evidence, E, can be taken to be something that she justifiably believes or knows, if she currently knows (or justifiably believes) a set of propositions E, she has very good reason to believe E at a later time. This is

because E entails itself and inferring what is entailed by my evidence is always a rational inference. There might be pragmatic reasons to not make trivial inferences⁵⁷, but these do not impugn the epistemic character of those inferences. Something therefore goes wrong, epistemically speaking when an agent stops believing E even though she has not acquired any defeaters for E. States in which an agent has forgotten evidence therefore are decidedly epistemically imperfect states which involve in one way or another, a failure to respond fully to one's reasons. This is doubly true of cases where an agent initially forms a belief on the basis of inadequate evidence and then at a later point forgets this evidence but retains her belief. The agent improperly responds to her reasons initially and then compounds this by losing access to those reasons while not adjusting her beliefs accordingly. However, as I have argued in Chapter 5, our concern is with whether people respond fully to the reasons they have. The conception of rationality we are working with is one according to which people respond fully to their reasons as a whole. If forgetting one's evidence involves a failure to respond to one's reasons, then forgotten evidence cases cannot be used to construct permissive cases. This would be because permissive cases need to be cases whereby two parties who share the same evidence are fully responding to their reasons but still disagree. Forgotten evidence cases therefore cannot be permissive cases. If so, then we can disregard alleged permissive cases based on forgotten evidence because they are not genuine permissive cases. Having dealt with McGrath's version of conservatism, I shall now proceed to address McCain's.

McCain's version of conservatism, which he calls properly formulated epistemic conservatism (PEC) can be defined in the following way:

“PEC: If S believes that p and p is not incoherent, then S is justified in retaining the belief that p and S remains justified in believing that p so long as p is not defeated for S.” (McCain 2008, p186)

McCain proposes two defeating conditions. The first Defeating Condition (DC1) is defined as follows:

“DC1: If S has better reasons for believing that $\sim P$ than S's reasons for believing that p, then S is no longer justified in believing that P.”
(ibid)

⁵⁷ For any belief P that we may have, we can always infer P or I am the king of the world. The inference may be extremely silly, but that would be for pragmatic reasons, not because the inference is itself epistemically defective.

The second defeating condition (DC2) is:

“DC2: If S has reasons for believing that $\sim P$ which are as good as S's reasons for believing that P and the belief that $\sim P$ coheres equally as well or better than the belief that P does with S's other beliefs, then S is no longer justified in believing that P.” (ibid)

McCain's defeating conditions are much stronger than McGrath's. In fact, McCain's defeating conditions are so strong, it is unclear if any permissive cases can be constructed. In order to construct a permissive case, there needs to be two people who have the same evidence, one of whom already happens to believe a proposition while the other does not. Conservatism is supposed to generate permissive cases by making the extra boost provided by believing a proposition make up for the shortfall in evidence. However, in this instance, if the evidence for a proposition P is insufficient to justify P to someone who does not already believe P, that must be because the reasons for P are just as good as or worse than the reasons for not-P. After all, if for S, the reasons for P were better than the reasons for not-P, then P would be justified for S even if S did not already believe that P. However, if the reasons for P are as good as or weaker than the reason for not-P, then at least one of DC1 or DC2 would be satisfied. In that case even for someone who already believes that P, P is not justified for them.

Perhaps this is an oversimplification. Suppose two agents S1 and S2 have identical evidence for P and their evidence for P is just as good as their evidence for not-P. However, for S1, P coheres with her non-evidential beliefs better than not-P. For S2, not-P coheres with her non-evidential beliefs better than P. Presumably then, since DC1 and DC2 are not triggered, P is justified for S1 and not-P is justified for S2. Supposedly, this might be a permissive case since different agents who share the same evidence are justified in holding different attitudes towards the same proposition. However, it is unclear whether this is genuinely a permissive case.

One reason to doubt that this is genuinely a permissive case is that there seem to be non-evidential beliefs. That is to say, there are beliefs which could cohere or conflict with P while not providing reasons to believe P or not-P. The only way such beliefs could exist but not generate reasons is, if these beliefs are themselves insufficiently supported by the evidence. After all, if they are supported by sufficient reasons, they would provide additional reasons to believe P or not-P and would therefore be evidence. If so, then S1 and S2 would not have the same evidence for

P. In order for them to have the same evidence, these other differing beliefs which could be relevant to whether P would have to be unjustified.

If these other beliefs were unjustified, then when S1 and S2 respond properly to their reasons, they will not retain these beliefs. If they drop these beliefs, then the only beliefs that they retain are the justified ones; i.e. the ones that make up their evidence. Given that they have identical evidence, and they do not differ in their non-evidential beliefs, S1 and S2 have no grounds to believe differently from each other. Therefore, this is not a permissive case. To sum up, it seems doubtful that McCain's defeating conditions permit the construction of permissive cases. Even cases which initially seem permissive are in fact really not, since, when people respond fully to their reasons, there are no beliefs apart from an agent's evidence which could cohere or conflict with the proposition in question.

Weakening Standard Conservatism by introducing more defeating conditions does not seem to be a promising strategy to construct permissive cases. The key difficulty seems to lie in the fact that it is hard to construct a permissive case which does not involve illicit bootstrapping or an illicit extra boost. Let us look at other ways of weakening Standard Conservatism to see if they are any more successful.

7.4.3 Mere Beliefs

A third way of weakening conservatism is to weaken the positive epistemic status conferred on a proposition by believing it. On Poston's (2012) view, conservatism applies only to what he calls mere belief (ibid: 520). A mere belief, by Poston's definition, is a belief about which there is no evidence, one way or another (ibid). For Poston, the conservative principle says that if an agent merely believes that P she justifiably believes that P (ibid). On this view, believing a proposition does not give it an extra boost over and above that provided by its evidence. Rather, when an agent encounters evidence regarding a proposition she initially merely believes, she subsequently no longer merely believes it since she has acquired evidence regarding that proposition. In that case, Poston's principle does not apply. Since it is not a mere belief, it is no longer justified simply in virtue of being believed⁵⁸. Poston's principle is sufficiently odd that I will split my discussion of it into two parts. The first part will address cases in which there is no evidence regarding the proposition and the second part will address cases in which there is at least some evidence regarding the proposition.

⁵⁸ Arguably an alternative way of characterising Poston's approach is as a way of narrowing the enabling conditions for a belief to be justified.

When there is no evidence regarding a proposition, Poston's principle is too strong a version of conservatism. Poston wishes to limit the application of conservatism to just certain framework beliefs (*ibid*), i.e. beliefs like 'induction is reliable' or that 'the external world exists'. However, his principle also permits certain obvious cases of illicit bootstrapping. Consider Foley's beach case. In that case, an agent happens to believe that there is an even number of grains of sand in a beach on the basis of no evidence at all. Since this is a mere belief, it is justified under Poston's principle. However, it is absurd that if someone arbitrarily forms a belief on the basis of no evidence at all, it is justified. Before having the belief, it is clearly irrational, even by Poston's standards, to form the belief. After, forming the belief, Poston's principle classifies the belief as rational. The principle thereby seems to license an illicit bootstrapping of the rational status of one's beliefs without a corresponding change in either one's belief or evidence.

The problem here is that Poston is unable to restrict the application of his principle to framework beliefs only. If he could restrict the application, we would be less willing to regard such cases as illicit bootstrapping. Even if he could restrict his principle to just framework beliefs, such beliefs are unlikely to give rise to permissive cases. It is doubtful that anyone outside of a philosophy classroom seriously denies that induction is reliable or that the external world exists. If so, then we cannot construct a permissive case since everyone shares these beliefs. To sum up, in no evidence situations, Poston's principle is either too strong or, if suitably circumscribed, incapable of generating permissive cases.

Let us now turn to cases in which there is at least some evidence. Of such cases, there are two things we might say. Firstly, we might say that no permissive cases could arise once there is at least some evidence regarding a proposition. To see why, consider two agents Sally and Susan. Sally has a mere belief that P, while Susan has a mere belief that not-P. According to Poston's principle, they can justifiably retain their respective beliefs only while they have no evidence regarding P. If they were to acquire some evidence regarding P, then they would not have mere beliefs anymore and hence, according to the principle, those beliefs would not be justified; at least not without the evidence itself supporting one or the other. Since the P and not-P are not merely believed by Sally and Susan respectively, there is no extra boost provided by their initial attitudes. Therefore, Sally is not entitled to form a higher confidence in P than Susan is. Hence, no permissive case can be constructed once we introduce some evidence about the proposition in question.

The second thing we might say is that Poston's principle which applies to mere beliefs is the closest we come to discussing ur-priors. Like mere-beliefs, ur-priors are those credences we can be presumed to form about various propositions when we have no evidence concerning them. This makes it tempting to extend Poston's principle to cover more fine grained attitudes.

Mere Attitude: For any set of doxastic attitudes that are merely held, so long as they are probabilistically coherent, they are justified.

Here I stipulate that to merely hold an attitude towards a proposition, the attitude towards the proposition must be held in the absence of any evidence about the proposition. However, Mere Attitude would not be of any help to permissive Bayesians. That would be because Bayesians update by conditionalizing on their priors. If Susan's prior probability for a proposition P is low while Sally's prior probability is high, Sally's posterior probability for P will be higher than Susan's, all else equal. On a Bayesian picture, an agent's priors continue to exert some effect on their posterior probabilities. Overwhelming amounts of evidence might be able to reduce the size of this effect, but will never be able to eliminate it entirely. However, on Poston's view, the introduction of evidence is supposed to eliminate any effect that the initial belief had because it ceases to be a mere belief. Similarly according to Mere Attitude, the initial attitudes that people have should have no effect on their final attitudes once they receive any amount of relevant evidence. There might be some other updating rule that could accommodate this principle, but if there is, it would not be permissive as people who have the same evidence could not rationally update to different attitudes just because they have different priors.

To sum up, Poston's principle is too strong in certain cases and licenses illicit bootstrapping. In other situations, it is too weak to generate permissive cases. In addition, the sharp discontinuity between how Poston's account of conservatism treats beliefs with and without evidence make it unsuitable to justify permissivism based on different ur-priors.

7.4.4 Dynamic Conservatism

Instead of defining conservatism as a principle which picks out certain propositions as being justified for an agent merely in virtue of the doxastic state she happens to be, epistemic conservatism can be expressed as a set of principles about belief revision. On Podgorski's account, which he calls dynamic conservatism, the

justification of any belief under consideration supervenes on the agent's evidence. More specifically, he accepts the following principle:

“Evidence Restriction (ER): The only considerations that, from the perspective of epistemic rationality, contribute to P's being worth believing are evidential— they bear positively on the belief's truth.”
(Podgorski 2016: 355)

ER seems antithetical to any version of conservatism. After all, ER seems to just be a re-statement of Evidentialism according to which:

Evidentialism: “Doxastic attitude D towards proposition P is epistemically justified for S at t if and only if having D towards P fits the evidence S has at t” (Feldman and Conee 1985: 15)

Podgorski denies this and distinguishes a belief being worth believing and it being rational. However, even if the two were conceptually distinct, it seems natural to suppose something like the following belief-worthiness-rationality bridge principle would be true.

State Oriented Worth Believing (SWB): “It is irrational to believe something that is not worth believing (from one's perspective).”(Podgorski 2016: 359)

SWB seems to be intuitively obvious. It is hard to see how it could ever be epistemically rational to believe something that is not worth believing. However, according to Podgorski, the reason why standard forms of conservatism are vulnerable to the bootstrapping objection is that epistemic conservatives tend to accept SWB but reject ER. According to Podgorski, this should be reversed. He claims that SWB is too strong and should be weakened further to the following:

Dynamic Worth Believing (DWB): it is irrational to conclude consideration of whether P by forming or reaffirming the belief that P if P is not worth believing (from one's perspective). (ibid: 363)

DWB is weaker than SWB in that it is entailed by SWB but does not in turn entail it. DWB opens the way for a version of conservatism by permitting agents to retain a belief P that is not worth believing so long as they have not considered whether P. This means that so long as an agent permissibly fails to consider whether P, she can permissibly believe that P even when the evidence does not support it. The

question then is: What are the circumstances in which an agent can permissibly not consider whether P? Podgorski supposes that there might be some circumstances in which we can permissibly not consider whether P even if we already believe it and our evidence is insufficient. He supposes the following:

“Inconsiderate: One is not always rationally required to initiate consideration whether P when one believes that P and one’s evidence does not make P worth believing (from one’s perspective).”
(ibid)

If Inconsiderate is true, someone may rationally acquire a belief, later receive disconfirming evidence but still maintain belief in P because she has permissibly not considered whether P since receiving the disconfirming evidence. This also allows us to see how a permissive case could be constructed. Two agents who possess the same evidence could rationally disagree about a proposition if it was worth considering for one of them but not the other. If such permissive cases can indeed be constructed, uniqueness is falsified.

I offer two arguments against this sort of permissive case. The first argument is that permissive cases cannot arise in political disagreement because in political disagreement, there are at least two conditions that make it worth considering your belief. The second argument is that the normative property we care about is the belief worthiness of a proposition, not necessarily its rationality in Podgorski’s sense of the term. I shall discuss each argument in turn.

In the context of politics, there are at least two conditions which can make it worth considering our beliefs. The first condition is that of acknowledged disagreement. The second condition is political or interpersonal justification. For the first condition, while I may not be required to consider whether my beliefs are true just because someone out there might disagree with me, this changes when I actually meet someone who does so. Consider the following principle:

Acknowledged disagreement: When we meet someone whom we believe disagrees with us about a proposition, we ought to consider whether our own beliefs are true.

When we see someone, who in the face of disagreement, does not even pause to consider whether they might be wrong, we are inclined to think that they are dogmatic and arrogant. This intuition supports the idea that permitting people to not

consider their beliefs in the face of disagreement would permit an unacceptable level of dogmatism.

However, it may seem that this requirement is too strong. Surely if I meet twenty people in succession, each five minutes apart and all of whom disagree with me, I am not obligated to reconsider my belief twenty times in the span of a few hours! After all, after re-examining my belief the first time, nothing could have changed after I last checked whether I had good enough reasons for my belief. I am unlikely to have gained any new evidence, apart from the bare fact of disagreement that bears on the issue. Moreover, surely I am not required to re-examine my beliefs right at the very instant I encounter disagreement. Intuitively, I should be allowed at least some time to re-examine my beliefs. We might fruitfully modify Acknowledged Disagreement without giving up entirely the requirement to re-examine one's beliefs in the face of disagreement. Firstly, we need to read Acknowledged Disagreement as permitting some, but not too long a time lag between meeting someone who we know disagrees with us and re-examining our beliefs. Secondly, we can exempt ourselves from examining our beliefs repeatedly if we meet someone else who disagrees with us before we even have a chance to re-examine our own beliefs after the first meeting. Thirdly, we can exempt ourselves from reconsidering our beliefs if we have already just done so thoroughly and the results of said re-examination are still fresh in our minds. With all these caveats and exemptions in place, Acknowledged Disagreement still has enough bite that when we meet others who disagree with us, we trigger re-consideration of our beliefs and end up believing according to the evidence.

It might be that if two people who disagree with each other live in their own corner of the world isolated from each other, they may be able to live their lives without having to reconsider their beliefs. However, in the context of democratic politics, the question of which candidate is the best or which policy is the best is front and centre. For instance, Susan and Sally each have to reconsider their own beliefs about abortion because they come to know of people who disagree with them about abortion. Suppose that a policy regulating abortion providers is on the ballot or up for vote in parliament. For Susan, who supports the policy, the mere fact that the regulation was previously not in place suggests that there is disagreement about the topic. After all, if everyone agreed that abortion ought to be regulated as per the policy, it is unlikely that it would only be going up for a vote now. Similarly, for Sally who opposes the policy, the fact that people are advancing a policy she opposes means that there are people who disagree with her on that topic. Given

Acknowledged Disagreement or any similar principle, Susan and Sally now have to reconsider their own beliefs about abortion.

Crucially, in a political liberal context the need to engage in political or interpersonal justification can also trigger considering our beliefs. The public justification requirement requires that laws be justifiable to all reasonable persons if they are to be legitimate. In practice, achieving this will require us to, firstly, consider whether we have good reasons for our policy preferences and, secondly, try to convince disagreeing others about the merits of our preferred laws and policies. Only when we discover that those who disagree with us have different and sufficient reasons of their own for their different policy preferences would it make sense for us to conclude that the policy cannot be justified to everyone. When deliberating about questions of policy or when attempting to justify or resist particular policy positions to others, we cite to ourselves or others our moral beliefs as reasons for our position. When we do that, we no longer merely passively believe those propositions, but are actively considering whether they are true. For instance, suppose I were to deliberate about whether to criminalise abortion, one reason I might invoke in favour is my belief that foetuses have a right to life. However, once I invoke that reason, I am not merely passively believing it, but actively considering it. Therefore, I am not permitted to continue believing that foetuses have the right to life if my evidence does not support it. To generalise, we might suppose that the following is a condition that justifies re-examining our beliefs.

Supporting Premise: If I ought to consider whether a given belief is true, I also ought to consider whether any of the premises employed in the relevant arguments for or against that belief are true.

If Supporting Premise is true, then, anytime someone claims to be able to reasonably reject a law because she accepts some conflicting belief, she ought to consider whether that belief is true. Similarly every time someone wishes to impose a law, she ought to consider whether the moral beliefs she invokes to justify the law are true. Insofar as both the person imposing the law and the person the law is being imposed upon ought to consider whether a given belief is true, that is not a permissive case.

To sum up, there seems to be at least two conditions in cases pertaining to political disagreement which make people's beliefs worth considering. Firstly, when people encounter disagreement with others during political campaigns, petitions and elections, it would be unacceptably dogmatic for them to continue on without ever

re-examining their beliefs to see if they fit the evidence. More importantly, when attempting to justify policies and laws to others or when attempting to justify why a proposed policy is unjust, people invoke their beliefs as sources of reasons. When doing so, this also triggers the need to re-examine those beliefs.

Notably, under Podgorski's account of conservatism, Susan and Sally in the Abortion case could not rationally disagree. Consequently, prior to coming to the university they may permissibly disagree since they may not be actively considering whether abortion is permissible. However, once at the university, when they encounter the various arguments for or against the permissibility of abortion, abortion not only becomes worth considering, Susan and Sally are actively considering whether abortion is permissible. Given Susan's and Sally's evidence, either it is worth believing that abortion is permissible, it is worth believing that abortion is not permissible or it is not worth believing either. Whichever of these three options turns out to be the case, Susan and Sally cannot rationally disagree. Any permissive case that could be based on Podgorski's account of conservatism has to involve at least one of the disagreeing parties permissibly retaining a belief against the evidence because she permissibly fails to consider whether her beliefs are true. However, no such case that would be relevant to the political context appears to be possible.

The second argument against such permissive cases rests on the observation that on Podgorski's view, there are two distinct normative questions that attach to a proposition. One question is whether a proposition is worth believing while the other is whether it is rational to believe it. We have hitherto assumed that rationality supervened on or was identical to belief-worthiness. However, given the rejection of SWB, there is a question as to which normative question we ought to care about *vis a vis* uniqueness.

Our concern with uniqueness arises from a concern about whether people who respond fully to their epistemic reasons can disagree given that they have the same total body of evidence. As I have argued in the fifth chapter, this is due to the fact that rationality is about responding perfectly to one's epistemic reasons. Inconsiderate permits people to sometimes not respond to their epistemic reasons. Notably, since, in Podgorski's view, people can be rational without necessarily responding to their epistemic reasons, Podgorski's sense of the term rationality is not the normative question we wish to answer. On the other hand, ER indicates that belief-worthiness supervenes on the agent's epistemic reasons. Therefore, given

that we are concerned about responsiveness to epistemic reasons, a situation counts as a permissive case only if between two people who have the same evidence, a proposition is worthy of belief for one of them but not for the other. However, there is no reason to think that any such case can be constructed.

It seems then that even on Podgorski's account of conservatism, we cannot construct a permissive case. As I have argued, other accounts of conservatism in the literature are either too strong, and therefore vulnerable to the bootstrapping objection, or they are too weak to construct a permissive case. If this is right, there is no plausible version of epistemic conservatism which would license permissivism.

With this, I have closed off the second way in which having different backgrounds could potentially be a source of rational disagreement without being a source of different evidence. Therefore, as far as the fourth burden of judgment is concerned, it cannot generate rational disagreement except by causing people to have different total bodies of evidence. In addition, by showing that people cannot have different epistemic standards arising from different initial beliefs, I have closed off the last way in which, as per the second burden of judgment, people can rationally disagree because they have different epistemic standards. If people cannot permissibly have different epistemic standards, then permissivism is false. The only way in which people could rationally disagree is if they had different total bodies of evidence.

Chapter 8 - Conclusion

To sum up, I have shown that while the standard account of political liberalism presupposes permissivism, permissivism is false. In the first chapter, I described what I take to be the standard account of political liberalism and showed how it related to reasonable disagreement and the burdens of judgment. I also argued that reasonable disagreement entailed rational disagreement. In the second chapter, I argued that the first, second, third, fifth and, to some extent, the fourth burdens of judgment presuppose epistemic permissivism.

In the subsequent five chapters, I proceeded to argue that permissivism is false by arguing for uniqueness. In my counterarguments in favour of the uniqueness thesis, in the third chapter, I have distinguished between intrapersonal and interpersonal uniqueness and argued that any plausible account of permissivism must be consistent with intrapersonal uniqueness. In the fourth chapter, I have argued that any account of permissivism which denies intrapersonal uniqueness about epistemic standards is subject to the same objections as versions of permissivism that deny intrapersonal uniqueness.

In the fifth through seventh chapters, I have presented three different indexing proposals for how an epistemic standard that is appropriate for an agent, could vary from agent to agent. I showed, for each indexing proposal, that no such account was successful. However, some way of indexing different epistemic standards to agents needs to be successful in order for versions of permissivism consistent with intrapersonal uniqueness about epistemic standards to be true. While the three accounts I present are by no means exhaustive of all such possible accounts, they cover the most obvious ones. Given these arguments, the burden is now on permissivists to account for how permissivism could be true. However, since it is unclear as to how such an argument would be shaped, we can be justified in being highly confident that permissivism is false.

What then follows from this? At the very minimum, the account of the burdens of judgment, that is, the causes of rational disagreement among reasonable persons requires extensive revision. Of the six burdens, the last burden is not a source of reasonable disagreement because it is ultimately not a source of disagreement at all. People, may, as a practical matter choose to make different trade-offs between their values even if they agreed on all their moral and religious doctrines. Of the remaining five, four of them presuppose permissivism which, as I have argued, is false. Of the remaining burden, the fourth burden, two of the three ways in which

people's different backgrounds could generate disagreement presuppose permissivism. The revised account of the burdens of judgment would contain only one item: people reasonably disagree because they have different bodies of evidence.

Presumably, this will severely limit the scope of reasonable disagreement. Far fewer claims are subject to reasonable disagreement. If the scope of reasonable disagreement is severely limited, much of the intuitive appeal of the political liberal project goes away. Intuitively, public reason is necessary because there is reasonable disagreement about a great many questions. If much fewer questions are subject to reasonable disagreement, then the perceived need for public reason is greatly reduced.

A contrary view is that apart from the specifics about the burdens of judgment, nothing much changes. While it is true that four of the six burdens of judgment presuppose epistemic permissivism, since not all of them do, the structure of political liberal theory would remain largely intact. The implicit claim is that reasonable persons can rationally disagree even if permissivism is false. The reason that they can do so is because reasonable persons can have different evidence about the propositions which justify laws and public policies. For instance, someone who has a degree in economics will possess evidence about the effects of wage or trade policies that someone who does not have said degree is not likely to have. Climate scientists will have evidence about global warming that lay persons may not have. We might generalise and say that for many empirical propositions which are relevant to public policy, people will have different evidence.

However, would the same thing be true of more basic moral propositions? If it can be shown that we share all the relevant evidence about some important basic moral doctrines, then, given that permissivism is false, reasonable disagreement about these doctrines is not possible. If so, then, at the very least, the core political liberal claim that certain illiberal laws are illegitimate because they are justified on the basis of moral or religious doctrines subject to reasonable disagreement is falsified. Relatedly, if empirical propositions are the only ones subject to reasonable disagreement, it does not seem illegitimate to impose a policy on people when they rationally disagree with the policy only because they lack certain pieces of empirical evidence. The intuition here is that only rational disagreement about basic moral

questions matters for legitimacy⁵⁹. If this is right, then the distinction between policies that are legitimate and those which are objectively correct seems to disappear. Working out whether basic moral propositions are subject to reasonable disagreement and if not, what the implications are for political liberalism is the subject of future work. One potentially fruitful line of enquiry is to attempt to formulate a version of political liberalism that does not rely on an account of reasonable disagreement.

⁵⁹ See for instance, Mill (1859) who argues that we may permissibly stop a man from crossing an unsafe bridge if the only reason he is trying to cross is that he does not know that the bridge is unsafe. The implication is that the public justification requirement for legitimacy applies primarily to normative disagreements. I take this to be the dominant view. For a contrary view, see Kappel (2017) who argues that imposing such policies on those who disagree with them because they lack access to the empirical evidence is illegitimate.

Bibliography

- Bajaj, Sameer (2017) Self-defeat and the Foundations of Public Reason, *Philosophical Studies* 174 (12): 3133–3151
- Ballantyne, Nathan and Coffman, E.J. (2011) Uniqueness, Evidence and Rationality, *Philosopher's Imprint*, 11(18): 1-13
- Berker, Selim (2013) The Rejection of Epistemic Consequentialism, *Philosophical Issues* 23: 363-387
- Boettcher, James W (2007) Respect, Recognition and Public Reason, *Social Theory and Practice* 33 (2): 223-249
- Burge, Tyler (1993) Content Preservation, *The Philosophical Review* 102(4): 457-488
- Card, David and Krueger, Alan B (1994) Minimum Wages and Employment: A Case-Study of the Fast Food Industry in New Jersey and Pennsylvania, *The American Economic Review* 84(4): 772-793
- Carnap, Rudolph (1950) *Logical Foundations of Probability*, University of Chicago Press
- Cassam, Qassim (2016) Vice Epistemology, *The Monist* 99: 159-180
- Chandler, Jacob (2014) Subjective Probabilities Need Not Be Sharp, *Erkenntnis* 79: 1273–1286.
- Chisholm, Roderick (1980) A Version of Foundationalism, *Midwest Studies in Philosophy* 5(4): 543-564
- Christensen, David (1994) Conservatism in Epistemology, *Nous* 28(1): 69-89
- Cohen, Stewart (1999) Contextualism, Skepticism and the Structure of Reasons, *Philosophical Perspectives* 13(Epistemology): 57-89
- Comesana, Juan (2010) Is Evidence Knowledge? *Philosophy and Phenomenological Research* 80(2): 447-454
- Conee, Earl and Feldman, Richard (2008) 15. Evidence in Smith, Quentin eds. *Epistemology: New Essays*, Oxford University Press

Corberos, Pablo (2011) Paraconsistent Vagueness: A Positive Argument, *Synthese* 183: 211-227

Cote-Bouchard, Charles (2015) Epistemic Instrumentalism and the Too Few Reasons Objection, *International Journal of Philosophical Studies* 23(3): 337-355

Davis, Wayne A. (2015) Knowledge Claims and Context: Belief, *Philosophical Studies* 172: 399-432

DeRose, Keith (1992) Contextualism and Knowledge Attributions, *Philosophy and Phenomenological Research*, 52(4): 913-929

Drake, Jonathan (2017) Doxastic Permissiveness and the Promise of Truth, *Synthese* 194: 4897-4912

Enoch, David (2017) Political Philosophy and Epistemology: The Case of Public Reason in Sobel, David, Vallantyne, Peter and Wall, Steven eds. *Oxford Studies in Political Philosophy* Vol 3: 132-163

Estlund, David (1998) The Insularity of the Reasonable: Why Political Liberalism must admit the Truth, *Ethics* 108(2): 252-275

Feldman, Richard and Conee, Earl (1985) Evidentialism, *Philosophical Studies* 48(1): 15-34

Feldman, Richard (1995) Authoritarian Epistemology, *Philosophical Topics* 23(1): 147-169

Feldman, Richard (2002) *Epistemology*, Prentice Hall

Fine, Kit (1975) Vagueness, Truth and Logic, *Synthese* 30(3/4): 265-300

Foley, Richard (1983) Epistemic Conservatism, *Philosophical Studies* 43(2): 165-182

Foley, Richard (1987) *The Theory of Epistemic Rationality*, Harvard University Press

Freeman, Samuel (2007), Rawls, Routledge

Gaus, Gerald (2011) *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*, Cambridge University Press

Goldman, Alvin I (2011) Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package in Dougherty, Trent ed. Evidentialism and its Discontents: 254-280

Greaves, Hilary and Wallace, David (2006) Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility, *Mind* 115(459): 607-632

Hambourger, Robert (1987) Justified Assertion and the Relativity of Knowledge, *Philosophical Studies* 51(2): 241-269

Harman, Gilbert (1973) *Thought*, Princeton University Press

Harman, Gilbert (1986) *Change in View*, The MIT Press

Hart, HLA (1961) *The Concept of Law*, Oxford University Press

Horowitz, Sophie (2018) Epistemic Value and the Jamesian Goals in Ahlstrom, H. Kristoffer and Dunn, Jeffrey eds. *Epistemic Consequentialism*, Oxford University Press

Howson, Colin (2003) Probability and Logic, *Journal of Applied Logic* 1: 151-165

Howson, Colin and Urbach, Peter (2005) *Scientific Reasoning: The Bayesian Approach* 3rd edition, Open Court

Howson, Colin (2007) Logic With Numbers, *Synthese* 156(3) Bayesian Epistemology: 491-512

Howson, Colin (2009) Can Logic be Combined with Probability? Probably, *Journal of Applied Logic* 7: 177-187

Huemer, Michael (1999) The Problem of Memory Knowledge, *Pacific Philosophical Quarterly* 80: 346-357

Hyde, Dominic and Colyvan, Mark (2008) Paraconsistent Vagueness: Why Not? *Australasian Journal of Logic* 6:107-121

James, William (1896) *The Will to Believe*, New World

Jeffreys, Harold (1948) *Theory of Probability*, Oxford University Press

Jonch-Clausen, Karin and Kappel, Klemens (2015) Social Epistemic Liberalism and the Problem of Deep Disagreements, *Ethical Theory and Moral Practice* 18: 371-384

Kappel, Klemens (2017) Fact-dependent Policy Disagreements and Political Legitimacy, *Ehtical Theory and Moral Practice* 20(2): 313-331

Kelly, Thomas (2003) Epistemic Rationality as Instrumental Rationality: A Critique, *Philosophy and Phenomenological Research* 66(3): 612-640

Kelly, Thomas (2010) Peer Disagreement and Higher Order Evidence in Feldman, Richard and Warfield, Ted A eds *Disagreement*, Oxford University Press: 111-174

Kelly, Thomas (2014) Evidence Can be Permissive in Steup, Mathias, Turri, John and Sosa, Ernest eds. *Contemporary Debates in Epistemology 2e*, Wiley-Blackwell: 298-312

Kopec, Matthew (2015) A Counterexample to the Uniqueness Thesis, *Philosophia* 43:403 – 409, DOI 10.1007/s11406-014-9579-x

Kopec, Matthew and Titelbaum, Michael G (2016) The Uniqueness Thesis, *Philosophy Compass* 11(4) : 189-200

Korcz, Keith (2000) The Causal-Doxastic Theory of the Basing Relation, *The Canadian Journal of Philosophy* 30(4): 525-550

Kolmogorov, Andrey (1956), *Foundations of the Theory of Probability*, Chelsea

Kornblith, Hilary (1983) Justified Belief and Epistemically Responsible action, *The Philosophical Review*, 92(1): 33-48

Kornblith, Hilary (1993) Epistemic Normativity, *Synthese* 94(3): 357-376

Kornblith, Hilary (2001) Epistemic Obligation and the Possibility of Internalism in Fairweather, Abrol and Zagzebski, Linda eds. *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Oxford University Press: 231-248

Kornblith, Hilary (2002) *Knowledge and its Place in Nature*, Oxford University Press

Kvanvig, Jonathan (2003) Propositionalism and the Perspictival Character of Justification, *American Philosophical Quarterly* 40(1): 3-18

Kvanvig, Jonathan (2014) *Rationality and Reflection*, Oxford University Press

Lackey, Jennifer (2010) What Should We Do When We Disagree? in: Gendler, Tamar Szabó and Hawthorne John eds. *Oxford Studies in Epistemology 3* Oxford University Press: 274-293.

- Larmore, Charles (2008) *The Autonomy of Morality* Cambridge University Press
- Larmore, Charles (2015) *Political Liberalism: Its Motivation and Goals* in Sobel, David, Vallentyne, Peter and Wall, Steven eds. *Oxford Studies in Political Philosophy* Vol 1 Oxford University Press: 63-88
- Laudan, Larry (1990a) Aimless epistemology, *Studies in History and Philosophy of Science Part A* 21(2): 315-322
- Laudan, Larry (1990b) Normative Naturalism, *Philosophy of Science* 57(1): 44-59
- Leite, Adam (2007) Epistemic Instrumentalism and Reasons for Belief, *Philosophy and Phenomenological Research* 75(2): 456-464
- Levinstein, Benjamin Anders (2017) Permissive Rationality and Sensitivity, *Philosophy and Phenomenological Research* 94(2): 342-370
- Lewis, David (1982) Logic for Equivocators, *Nous* 16(3): 431-441
- Lewis, David (1986) *A Subjectivist's Guide to Objective Chance* in Lewis, David ed. *Philosophical Papers* Vol 2, Oxford University Press
- Lister, Andrew (2016) Public Reason, Liberal Neutrality and (Same Sex) Marriage in den Otter, Ronald ed *Beyond Same-Sex Marriage*, Lexington Books
- Lockard, Matthew (2013) Epistemic Instrumentalism *Synthese* 190: 1701-1718
- Lycan, William (1988) *Judgment and Justification*, Cambridge University Press
- Maffie, James (1990) Naturalism and the Normativity of Epistemology, *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 59(3): 333-349
- Matheson, Jonathan (2018) Disagreement and Epistemic Peers, in *Oxford Handbooks Online* DOI: 10.1093/oxfordhb/9780199935314.013.13
- May, Simon Cabulea (2016) Liberal Neutrality and Civil Marriage in Brake, E ed *After Marriage*, Oxford University Press: 9-28
- McCain, Kevin (2008) The Virtues of Epistemic Conservatism, *Synthese* 164: 185-200
- McGrath, Matthew (2007) Memory and Epistemic Conservatism, *Synthese* 157: 1-24

- Meacham, Christopher J. G. (2014) Impermissive Bayesianism, *Erkenn* 79: 1185-1217
- Mill, John Stuart (1859) *On Liberty*, reprinted by Batoche Books, Kitchener (2001)
- Neumark, David and Wascher, William (1995) The Effects of Minimum Wages on Teenage Employment and Enrollment: Estimates from Matched CPS Data, NBER Working Papers 5092, National Bureau of Economic Research, Inc.
- Nozick, Robert (1993) *The Nature of Rationality*, Princeton University Press
- Papineau, David (1999) Normativity and Judgment, *Aristotelian Society Supplementary Volume* 73(1): 17-43
- Parfit, Derek (2011) *On What Matters Vol 1*, Oxford University Press
- Peter, Fabienne (2018) Political Legitimacy under Epistemic Constraints: Why Public Reasons Matter in Schwartzberg, Melissa and Knight, Jack eds. *NOMOS LXI*
- Pettigrew, Richard (2016) Jamesian Epistemology Formalised: An explication of The Will to Believe, *Episteme* 13(3): 253-268
- Podgorski, Abelard (2016) Dynamic Conservatism, *Ergo* 3(13): 349-376
- Pollock, John and Cruz, Joseph (1999) *Contemporary Theories of Knowledge* 2nd ed, Rowman and Littlefield
- Poston, Ted (2012) Is there an I in Epistemology, *Dialectica* 66(4): 517-541
- Quine, Willard Van Orman (1963) Two Dogmas of Empiricism in Quine W.V.O eds *From a Logical Point of View*, Harper Torch Books
- Quong, Jonathan (2010) *Liberalism Without Perfection*, Oxford University Press
- Ramsey, F. P. (1931) Truth and Probability in Braithwaite, R. B. eds. *Foundations of Mathematics and other Essays*, London: Kegan, Paul, Trench, Trubner, & Co: 156–198
- Rawls, John (1993) *Political Liberalism*, Columbia University Press
- Rawls, John (2001) *Justice as Fairness: A Restatement*, Harvard University Press
- Rosa, Luis (2012) Justification and The Uniqueness Thesis, *Logos and Episteme* 3(4): 571-577

Rosa, Luis (2016) Justification and The Uniqueness Thesis Again - A Response to Anantharaman, *Logos and Episteme* 7(1): 91-100

Scanlon, T.M. (2013) *Being Realistic About Reasons*, Oxford University Press

Schoenfield, Miriam (2014) Permission to Believe: Why Permissivism is True and What it Tells us on Irrelevant Influences on Belief, *Nous* 48(2): 193-218

Sklar, Lawrence (1975) Methodological Conservatism, *The Philosophical Review* 84(3): 374-400

Simpson, Robert Mark (2017) Permissivism and the Arbitrariness Objection, *Episteme* 14(4): 519-538

Swain, Marshall (1979) Justification and the Basis of Belief in Pappas G.S. ed. *Justification and Knowledge*, Philosophical Studies Series in Philosophy vol 17, Springer, Dordrecht

Thagard, Paul (1982) From the Descriptive to the Normative in Psychology and Logic, *Philosophy of Science* 49(1): 24-42

Turri, John (2010) On the Relationship between Propositional and Doxastic Justification, *Philosophy and Phenomenological Research* 80(2): 312-326

Vahid, Hamid (2004) Varieties of Epistemic Conservatism, *Synthese*, 141(1): 97-122

Vallier, Kevin (2011) Convergence and Consensus in Public Reason, *Public Affairs Quarterly*, 25(4): 261-279

van Frassen, Bas (1990) Figures in a Probability Landscape in Dunn, J. Michael and Gupta, Anil eds. *Truth or Consequences: Essays in Honor of Nuel Belnap*, Kluwer Academic Publishers, Dordrecht: 345-356

van Inwagen, Peter (2010) *We're Right, They're Wrong* in Feldman, Richard and Warfield, Ted A eds. *Disagreement*, Oxford University Press: 10-28

Wall, Steven (2016) The Pure Theory of Public Justification, *Social Philosophy and Policy* 33(2): 204-226

Wedgwood, Ralph (2008) Contextualism about justified Belief, *Philosopher's Imprint* 8(9): 1-20

White, Roger (2005) Epistemic Permissiveness, *Philosophical Perspectives* 19(1): 445- 459

White, Roger (2010) Evidential Symmetry and Mushy Credence, in Gendler, Tamar Szabo and Hawthorn, John eds. *Oxford Studies in Epistemology* 3: 161-186.

White, Roger (2014) Evidence Cannot be Permissive in Steup, Mathias, Turri, John and Sosa, Ernest eds. *Contemporary Debates in Epistemology* 2e, Wiley-Blackwell: 312-322

Williamson, Timothy (1994) *Vagueness*, Routledge

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