

Christopher Lawrence

Permission to believe:

Descriptive and prescriptive beliefs in the
Clifford/James debate

Thesis Presented for the Degree of
DOCTOR OF PHILOSOPHY
in the Department of Philosophy
UNIVERSITY OF CAPE TOWN

August 2020

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Acknowledgements

With thanks to my supervisor Dr Jack Ritchie, but most of all to my wife Wendy for putting up with me over all the highs and lows.

This thesis represents my own work, in both concept and execution.

Signed by candidate

Christopher Lawrence

Abstract

William Clifford's 'The Ethics of Belief' proposes an 'evidence principle':

...it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence (1877, 1879:186).

Its universal, absolutist language seems to hide something fundamentally correct.

We first argue for excluding prescriptive beliefs, and then consider further apparent counter-examples, culminating in more restricted, qualified wording:

If anything is morally wrong, then it is morally wrong within the category of descriptive belief to believe anything knowingly or irresponsibly on insufficient evidence in the absence of any conflicting and overriding moral imperative except when the unjustified believing is outside the believer's voluntary control.

We test this against William James's counter-claim for qualified legitimate over-belief ('The Will To Believe', 1896, 2000), and suggest additional benefits of adopting an evidence principle in relation to the structured combinations of descriptive and prescriptive components common to religious belief.

In search of criteria for 'sufficient' and 'insufficient' evidence we then consider an 'enriched' Bayesianism within normative decision theory, which helps explain good doxastic practice under risk. 'Lottery paradox' cases however undermine the idea of an evidence threshold: we would say we justifiably believe one hypothesis while saying another, at the same credence level, is only very probably true.

We consider approaches to 'pragmatic encroachment', suggesting a parallel between 'practical interest' and the 'personal utility' denominating the stakes of the imaginary gambles which Bayesian credences can be illustrated as. But personal utility seems inappropriately agent-relative for a moral principle.

We return to Clifford's conception of our shared responsibilities to our shared epistemic asset. This 'practical interest we ought to have' offers an explanation for our duty, as members of an epistemic community, to get and evaluate evidence; and for the 'utility' stakes of Bayesian imaginary gambles. Helped by Edward Craig's (1990, 1999) 'state-of-nature' theory of knowledge it provides a minimum threshold to avoid insufficient evidence and suggests an aspirational criterion of sufficient evidence:

Wherever possible, a level of evidence sufficient to support the level of justification required to be a good informant, whatever the particular circumstances of the inquirer.

Table of Contents

Acknowledgements	2
Abstract	4
Summary	9
Chapter 1 Introduction	13
1.1 William Clifford: ‘The Ethics of Belief’	13
Sins of belief	14
Voluntary and involuntary beliefs	17
Rationale behind CP	18
1.2 Drafting the principle	21
Relation to other moral principles	21
What is meant by ‘belief’?	22
An everyday example	23
1.3 Clarifying the principle	24
Jonathan Adler: <i>Belief’s Own Ethics</i>	24
Susan Haack: <i>“The Ethics of Belief” Reconsidered</i>	31
Social context of belief: 1	35
Why both?	38
1.4 Reasons for holding an evidence principle	41
1.5 Conclusions	42
Chapter 2 Prescriptive beliefs	43
2.1 Introduction	43
2.2 Hypothetical prescriptive beliefs	44
2.3 Evidence and truth value	46
Reasons for a prescriptive belief	47
Truth value of a prescriptive belief	49
2.4 Clifford on moral beliefs	52
2.5 William James on prescriptive beliefs	57
2.6 Problems with Clifford’s principle applied to prescriptive beliefs	62
First level	63
Second level	64
Third level	68
Fourth level	72
Fifth level	72
2.7 Conclusions	73
Chapter 3 Against an evidence principle	74
3.1 Involuntary beliefs	74
3.2 Foundational beliefs	81
3.3 Past and future; testimony; uniformity of nature	90
Beliefs about the past	91
Testimony	91
Uniformity of nature	95
Beliefs about the future	96

Belief as a guide to action	99
Belief as intrinsically forward-looking	102
Conclusion.....	109
3.4 Personal relations and self-fulfilling beliefs.....	110
Faith in a fact.....	110
Smith and Jones	116
Husband and wife.....	117
3.5 Personal preferences.....	119
3.6 Aesthetic beliefs	120
3.7 Apparently harmless over-beliefs.....	121
3.8 Private over-beliefs	122
3.9 Unjustified true beliefs.....	125
3.10 Conclusions.....	126
Chapter 4 William James and the believing attitude	129
4.1 Introduction	129
4.2 Compound structure of religious belief	131
EP4 and prescriptive beliefs.....	131
Descriptive and prescriptive components	131
4.3 James’s belief framework.....	140
Contrast with Clifford	140
Genuine option	142
Sequence of believing.....	144
Pascal’s Wager	145
James on Pascal.....	147
Belief and the will.....	156
Our passional nature	160
Empiricism, absolutism and objective evidence	161
Knowing the truth and avoiding error	163
4.4 James’s case for religious belief.....	165
James’s religious hypothesis.....	166
Religious belief as a genuine option.....	167
Person to person	179
At our own risk	181
Moral sovereignty	185
Sovereignty dilemma revisited.....	190
Alternatives to an evidence principle.....	195
4.5 Conclusions	201
Summary of previous section.....	201
Broader context and chapter conclusions	203
Chapter 5 Sufficient and insufficient evidence	205
5.1 Introduction	205
5.2 How a criterion of sufficient evidence relates to EP4	208
5.3 Social context of belief: 2	210
5.4 Bayesian approaches.....	213
Bayes’ Rule applied to a test case	216
Evidence and normativity	226
Evidence threshold	230

Lottery paradox	233
5.5 Pragmatic encroachment.....	240
Stanley: What is at stake.....	240
Clarke: Degrees of belief vary by context	244
Practical interests	246
Whose practical interest? Whose personal utility?	249
5.6 Indefinite intersubjectivity	253
Shared facts and the ethics of belief	253
Craig’s ‘good informant’	255
A criterion of sufficient evidence for EP4	264
Concluding summary	267
Chapter 6 In conclusion	272
6.2 Summary	272
6.2 Status of the principle	274
Appendix A1 Clifford on the evolution of conscience	276
Appendix A2 Bayes’ Rule, evidence and prescriptive beliefs	281
Appendix A3 EP4 and subjective Bayesian permissiveness	285
Appendix A4 Bertrand’s paradox.....	289
A4.1 Three possible answers	289
A4.1.1 Random points on the circumference	289
A4.1.2 Randomly centred on randomly selected diameter	290
A4.1.3 Randomly selected midpoint	290
A4.2 Discussion	291
Appendix A5 Future directions	292
A5.1 Transcendental argument.....	292
A5.2 Scientific revolutions	293
A5.3 Moral revolutions	294
A5.4 Moral encroachment	295
Bibliography	296

Summary

In 'The Ethics of Belief' William Clifford proposes an 'evidence principle':

...it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence (1877, 1879:186)

The universal, absolutist terms in which it is expressed expose it to attacks which a more restricted, qualified principle might resist. This dissertation was therefore prompted by the suspicion that, despite its failings, Clifford's principle had got something fundamentally right. Its aim was to tease out what that something might be. For example it questions whether a principle like Clifford's can usefully apply to prescriptive beliefs, and concludes that attempting to do so seems fraught with obstacles, for little positive gain.

Chapter 1 introduces Clifford's principle and articulates his rationale for it. It then makes a prima facie case for an evidence principle less universally worded than Clifford's. Like Clifford's however our modified principle applies primarily to believing behaviour, and only secondarily to beliefs themselves.

In **Chapter 2** we give two reasons for excluding prescriptive beliefs. One is the challenge of explaining how evidence can support purely prescriptive beliefs. The other is that applying the principle to prescriptive beliefs seems incoherent – and not just because, as holding the principle would amount to holding a prescriptive belief, an evidence principle would have to presuppose itself in order to obey itself.

Chapter 3 explores the idea of an evidence principle restricted to descriptive beliefs and considers apparent counter-examples. These lead to clarifications and qualifications, culminating in our final wording:

If anything is morally wrong, then it is morally wrong within the category of descriptive belief to believe anything knowingly or irresponsibly on insufficient evidence in the absence of any conflicting and overriding moral imperative except when the unjustified believing is outside the believer's

voluntary control.

This makes explicit what Clifford left implicit: that it is a *morally* prescriptive principle.

Chapter 4 then tests our principle against William James's very different focus on the conditions for legitimate over-belief in general and religious over-belief in particular (in 'The Will To Believe', 1896, 2000). We also suggest a number of additional potential benefits, specifically in relation to religious belief, of adopting an evidence principle. These benefits all relate to the compound nature of much religious belief, consisting of structured combinations of descriptive and prescriptive components.

Contrasting moral positions are however notoriously hard to secure or demolish. Complicating the case of James versus Clifford is a clash between two different existential perspectives and orientations. James prioritises individual liberty, whereas Clifford and our own principle both foreground the social and communal consequences of epistemic behaviour. If our principle does not conquer, it seems at least to survive unscathed.

The final step, in **Chapter 5**, is to articulate a criterion of 'sufficient' evidence (and therefore of 'insufficient' evidence) which our evidence principle would require. The concept of the shared epistemic asset underpinning Clifford's principle ends up doing much of the heavy lifting in relation to our own. But first we see how far we get by considering Bayesian accounts of the relation between evidence and belief.

Bayesianism sees degree of belief as a personal probability which can be modelled as the betting rate an individual would accept as fair in an imaginary gamble. An initial challenge is the 'problem of the priors', to which our response is to consider only credences after at least one update cycle. We therefore assume a duty to get evidence and to take that evidence seriously. This makes sense if we adopt an 'enriched' Bayesianism considered within normative decision theory, which also explains both the normative aspect of Bayesianism ('why be a

Bayesian?') and the currency of personal utility which Bayesian imaginary gambles could be denominated in.

Although evidential sufficiency and insufficiency are not strictly part of Bayesianism, this does not make Bayesianism irrelevant to our principle. Bayesianism goes a long way towards explaining 'good practice' in our doxastic behaviour in support of decision-making under risk, particularly in respect of the relation between degrees of belief and evidential support and the avoidance of doxastic fallacies.

A key challenge though is to define sufficient evidence from a Bayesian foundation. We need a criterion of sufficient evidence short of conclusive evidence, as conclusive evidence needs no Bayesian explanation.

We consider the idea of a threshold. But apart from the issue of deciding and agreeing its numerical level, the mere idea of a threshold seems unworkable on its own. 'Lottery paradox' cases show we could have two beliefs at the same personal probability, one we would say we justifiably believe (or even know), and the other we would say is only very probably true.

To address this we consider two different approaches to 'pragmatic encroachment'. This seems initially promising, and in line with our assumed 'enriched' Bayesianism embedded in normative decision theory: the 'practical interest' motivating pragmatic encroachment seeming equivalent to the 'personal utility' representing the stakes of the imaginary gambles which Bayesian credences can be illustrated as.

But a conception of practical interest restricted to personal utility seems inappropriately agent-relative for an evidence principle expressed in moral terms. This is particularly evident when relating our developing notion of evidential sufficiency to, for example, Clifford's ship owner story. From the perspective of personal utility, the ship owner's interest was served by believing as he did, not by withholding belief on insufficient evidence.

For our purposes therefore there seems every reason to incorporate Clifford's dual conception of our shared epistemic asset and our shared responsibilities to that asset. This transforms the 'practical interest' in pragmatic encroachment into the 'practical interest we ought to have', which then provides an explanation for our duty to get evidence and take evidence seriously, and for the 'utility' currency denominating the stakes of Bayesian imaginary gambles. It also removes the problematic agent-relativity, and allows us to formulate a criterion of evidential sufficiency for belief, with the help of Edward Craig's (1990, 1999) 'state-of-nature' theory of knowledge:

Wherever possible, a level of evidence sufficient to support the level of justification required to be a good informant, whatever the particular circumstances of the inquirer.

The imperative to be a 'good informant' commits us to intersubjectivity, allowing us to navigate around radical scepticism and avoid consigning virtually all our beliefs to a black hole of solipsism.

Chapter 6 then summarises what conclusions the overall dissertation can support, and examines what the status of our principle might be.

Chapter 1 | Introduction

This chapter introduces William Clifford's evidence principle,¹ which we label 'CP' ('Clifford's Principle'). We then make an initial prima facie case for a modified principle based on CP, but restricted to descriptive beliefs. Like CP however our principle applies primarily to believing behaviour and only secondarily to the resultant beliefs themselves. To start to assess the universality of our principle, we pick an everyday descriptive belief and ask if our ordinary expectations regarding belief language would make a believer culpable if he² acquired, held or asserted the belief without sufficient evidence. We sample recent literature to explore whether obligation is implicit in the conceptual structure of belief; epistemic versus moral appraisal of belief; and belief in its social context.

1.1 William Clifford: 'The Ethics of Belief'

In 'The Ethics of Belief' Clifford defends his evidence principle:

CP *...it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence* (1877, 1879:186).

Below we unpack Clifford's apparent intentions: that CP applies to all beliefs, including prescriptive beliefs; that 'wrong' means 'morally wrong';³ that CP applies primarily to believing behaviour; and that 'always, everywhere and for anyone' is meant literally. We also articulate what Clifford seems to take as true about people and their world, to warrant a principle like CP.

¹ Clifford's original principle, and any appropriate modification of it, will be referred to as 'evidence' principles rather than 'evidentialist' principles. This is to avoid the impression that the principles exactly correspond to any particular flavour of 'evidentialism' discussed in the literature, for example Adler's 'intrinsic' evidentialism (2002:2), which he distinguishes from both 'moderate' evidentialism ('Evidence should generally determine the strength of belief, but not always' (2002:3)) and the 'traditional' evidentialism of, say, Locke or Hume ('the strength of one's belief ought to be proportional to the strength of one's reasons' (2002:25)).

² Alternate chapters will default to masculine and feminine pronouns.

³ In much of the literature 'ethics' in 'ethics of belief' is not synonymous with 'morality', as the 'ethics of belief' could consider prudential or moral or epistemic norms, separately or in combination. See for example Chignell (2018: §2.1). For clarity we will generally use 'moral' and its cognates rather than 'ethical' when a specifically moral context is intended.

Clifford delineates what we may believe (positive permission) from what we may not believe (negative obligation), saying little about what we must believe (positive obligation) or do not have to believe (negative permission). We will also prioritise those first two categories.

Sins of belief

Clifford applies CP to both descriptive and prescriptive beliefs. For our purposes a descriptive belief is about what is (or is not) the case. A prescriptive (or normative⁴) belief is about what ought (or ought not) to be the case. Prescriptive beliefs include moral beliefs, but not all prescriptive beliefs are moral beliefs. Adopting CP would itself count as adopting a morally prescriptive belief alongside, say, 'cheating is wrong' and 'thou shalt not kill'. We are not assuming all beliefs must be either descriptive or prescriptive. Aesthetic beliefs for example may qualify as neither.⁵

Clifford applied CP to the prescriptive belief that it is wrong to 'steal and tell lies' (1877, 1879:188). However his keynote story of the passenger ship owner (1877, 1879:177) features a purely descriptive belief. The owner suspects his ship might be unseaworthy, but manages to overcome his doubts: not by having her overhauled and refitted but by trusting in Providence. She sails and sinks in mid-ocean. Is he guilty of the death of passengers and crew? Undoubtedly. He 'did sincerely believe in the soundness of his ship', but he had 'acquired his belief not by honestly earning it in patient investigation, but by stifling his doubts' (1877, 1879:178).

What if the ship had been sound all along? Clifford claims he would still have been guilty. Not of manslaughter of course, but of holding a belief he had no right to hold, its origin being faulty. He would still have endangered the lives on board.⁶

⁴ We will use either depending on context. No significant distinction is intended.

⁵ See also p120.

⁶ See also p125ff.

Clifford argues that guilt attaches to the unjustified belief ('over-belief'⁷) itself, not just to action arising from it. This is because 'it is not possible so to sever the belief from the action it suggests as to condemn the one without condemning the other' (1877, 1879:181). No 'real belief, however trifling and fragmentary it may seem, is ever truly insignificant' (1877, 1879:181) or without effect in the world, if only because each over-belief makes the next one easier to acquire.⁸ Over-beliefs are insidiously corrupting because each

prepares us to receive more of its like, confirms those which resembled it before, and weakens others; and so gradually it lays a stealthy train in our inmost thoughts, which may someday explode into overt action, and leave its stamp upon our character for ever. (1877, 1879:181-2)

Clifford thinks it not just unsound, unwise or an epistemic failing but morally wrong to over-believe:

*We all suffer severely enough from the maintenance and support of false beliefs and the **fatally wrong actions** which they lead to, and the **evil** born when one such belief is entertained is great and wide. But a **greater and wider evil** arises when the credulous character is maintained and supported, when a habit of believing for **unworthy reasons** is fostered and made permanent.* [Emphases added.] (1877, 1879:185)

As with other moral failings both immediate and indirect consequences may ensue. If I steal money, the theft itself may be insignificant: my victim 'may not feel the loss, or it may prevent him from using the money badly'. What 'hurts society' though is not so much the loss of property but the risk of becoming 'a den of thieves, for then it must cease to be society'. Similarly, over-belief is wrong primarily because of its effects on the community. The direct impact of my individual over-belief may be insignificant, but

⁷ See James (1907, 2000:131).

⁸ Examples of apparently harmless over-beliefs are discussed later: p121ff.

I cannot help doing this great wrong towards Man, that I make myself credulous. The danger to society is not merely that it should believe wrong things ... but that it should become credulous, and lose the habit of testing things and inquiring into them; for then it must sink back into savagery. [Emphases added.] (1877, 1879:185-6)

An individual's credulity is no 'private matter' as belief is a shared social asset:

Our lives are guided by that general conception of the course of things which has been created by society for social purposes. Our words, our phrases, our forms and processes and modes of thought, are common property, ... an heirloom which every succeeding generation inherits as a precious deposit and a sacred trust to be handed on to the next one, not unchanged but enlarged and purified... Into this, for good or ill, is woven every belief of every man who has speech of his fellows. (1877, 1879:182)

Because the wrong of over-belief is to society at large, Clifford sees us all obligated regardless of rank or role:

... Every rustic who delivers in the village alehouse his slow, infrequent sentences, may help to kill or keep alive the fatal superstitions which clog his race. Every hard-worked wife of an artisan may transmit to her children beliefs which shall knit society together, or rend it in pieces. (1877, 1879:183)

We all share a moral duty to avoid over-belief and so prevent intellectually stultifying credulity. We should all therefore foster good habits of belief and avoid bad habits:

Habitual want of care about what I believe leads to habitual want of care in others about the truth of what is told to me. (1877, 1879:186)

Voluntary and involuntary beliefs

Clifford includes voluntary and involuntary beliefs. Taken literally CP covers 'all beliefs', regardless of how much volition is involved. But his presentation confers praise and blame paradigmatically on those aspects of belief-acquisition and belief-rejection which are (or could be, or should be) under voluntary control. Hence his claim that unjustified belief is insidiously corrupting: a 'pestilence' (1877, 1879:184). Within each individual, if we

let ourselves believe for unworthy reasons, we weaken our powers of self-control, of doubting, of judicially and fairly weighing evidence. (1877, 1879:185)

At the extreme our self-control and discernment could so atrophy that our gullibility is completely involuntary, until in theory we could claim innocence for any resultant over-belief. But we would be wrong to let ourselves get to this state. Compare a father so tired from partying every night that he loses patience with his toddler son over something trivial and causes an injury. Perhaps the more extreme the exhaustion the less responsible the father was for the injury itself, but the more culpable he is for letting exhaustion deplete his self-control. CP can similarly cover involuntarily acquired beliefs, if the methods and habits cultivated for belief acquisition⁹ result from voluntary choices and omissions.

Clifford recognises cases of conscious volition exercised over beliefs. The ship-owner 'succeeded in overcoming [his] melancholy reflections' on the state of his ship and 'dismiss[ed] from his mind all ungenerous suspicions about the honesty of builders and contractors'. He 'acquired his belief ... by stifling his doubts', so 'inasmuch as he had knowingly and willingly worked himself into that frame of mind, he must be held responsible for it' (1877, 1879:177-8).

An example of less direct control is the 'agitating society' accusing members of an unorthodox religion of misdemeanours. Yet 'the evidence of their innocence was such as the agitators might easily have obtained, if they had attempted a fair

⁹ We will refer to these as 'doxastic' methods and habits.

inquiry' (1877, 1879:179). The agitators' beliefs resulted from their deliberate choices and omissions, but they did not necessarily intend having the beliefs they ended up holding.

Still further towards the involuntary extreme is the man whose 'belief is so fixed that he cannot think otherwise'. But even he

still has a choice in regard to the action suggested by it, and so cannot escape the duty of investigating on the ground of the strength of his convictions (1877, 1879:180-1)

Clifford interweaves voluntary and involuntary aspects of belief, within one individual and between individuals in the same community. Cultivating poor doxastic¹⁰ habits by voluntary choice and omission leads individuals and communities to 'lose the habit of testing things and inquiring into them', becoming more credulous and prone to wrongly formed involuntary beliefs.

Rationale behind CP

Clifford's case for CP appears to incorporate or presuppose a number of descriptive, largely empirical, claims, not all spelled out. Below is an attempt to articulate them:

- (1.1.1) Our beliefs generally influence our actions.
- (1.1.2) Our actions generally have consequences for others.
- (1.1.3) We have limited control over the degree to which, and manner in which, our beliefs influence our actions.
- (1.1.4) We make inferences about people's beliefs from their actions.
- (1.1.5) Beliefs supported by sufficient evidence are generally speaking more

¹⁰ See previous footnote.

likely to be true than those not so supported.

- (1.1.6) We generally take people's beliefs to be supported by sufficient evidence, and therefore likely to be true.
- (1.1.7) Voluntary behaviour in relation to how conscientiously we ensure our beliefs are supported by evidence significantly influences our doxastic habits, and therefore our future voluntary and/or involuntary beliefs.
- (1.1.8) We have limited control over how our voluntary behaviour influences our doxastic habits: see (1.1.7).
- (1.1.9) As inter-communicating social beings our individual beliefs and doxastic habits inter-relate and exert a powerful aggregate influence at community level. This aggregate of beliefs and doxastic habits represents a shared social asset inherited from one generation to the next.
- (1.1.10) We have limited control over how our individual beliefs and doxastic habits impact that shared social asset (1.1.9).
- (1.1.11) Our survival and well-being, at both individual and community level, are generally speaking maximised when future states of affairs correspond with our expectations, including expectations about the results of our deliberate actions.
- (1.1.12) Our survival and well-being, at both individual and community level, are generally speaking maximised when beliefs guiding our expectations and deliberate actions are supported by sufficient evidence and therefore most likely to be true. The shared social asset (1.1.9) is accordingly most effective in promoting community survival, well-being and cohesion when its component beliefs are supported by sufficient evidence and therefore most likely to be true, and when our shared doxastic habits encourage the acquisition of such beliefs.

(1.1.13) Generally speaking we each do what we can to maximise our own individual survival and well-being.

His case also incorporates or presupposes a generic prescriptive moral claim like:

(1.2.1) Other things being equal we should each do what we can to maximise the survival and well-being of others, at both individual and community level.

From (1.2.1) and the set of descriptive claims (1.1.1–13) can be derived a more specific prescriptive moral claim like:

(1.2.2) Other things being equal we should each do what we can to develop and encourage doxastic habits in ourselves and others aimed at ensuring our beliefs are adequately supported by evidence.

CP re-expresses the ‘positive obligation’ of (1.2.2) in terms of ‘negative obligation’.¹¹

Empirical (descriptive) claims can support moral (prescriptive) beliefs as long as there is at least one appropriately worded prescriptive statement – here (1.2.1) – to effect the is-to-ought transition.¹² ‘Thou shalt not kill’ for example appears to presuppose the empirical truth that ‘humans are mortal’, which on its own does not entail the sixth commandment.¹³

The empirical claims (1.1.1–13) appear at least plausible: if they were not broadly true a principle like CP would have little point. This dissertation will not investigate whether (1.1.1–13) are empirically true, but to what extent a principle like CP is conceptually sound. In general we will assume the psychological and sociological facts are broadly as in (1.1.1–13); and presuppose an overall meta-ethical position

¹¹ See p13.

¹² See p134, where a similar point is discussed in a different context.

¹³ Or count as evidence for it: see **Chapter 2**, particularly pp62–73.

incorporating a generic moral imperative like (1.2.1). The claim is therefore that someone who considers (1.1.1–13) substantially true and holds a generic moral belief like (1.2.1) should for consistency adopt a principle like CP.

CP has been critiqued from various perspectives, including whether it involves what today would be called ‘voluntarism’, and therefore fails if voluntarism fails.¹⁴ Future chapters will consider a range of objections, but in relation to a principle less universally worded than Clifford’s. The next section starts to sketch a principle which preserves the spirit of CP, shares the rationale articulated in (1.1.1–13) and (1.2.1), but should prove easier to defend.

1.2 Drafting the principle

Unlike CP our principle will only cover descriptive beliefs. But like CP it will be a moral principle, typically applying to descriptive beliefs acquired and/or held in contexts involving a degree of volition. Other things being equal the believer would be culpable insofar as he acquired and/or held the descriptive over-belief as a result of something he voluntarily did or refrained from doing.¹⁵ Like CP therefore it applies primarily to doxastic behaviour and only secondarily to the resultant beliefs.

Relation to other moral principles

The claim that over-belief (over-believing) in general, or any specific category of it, is morally wrong parallels the claim that (say) lying or deceiving are morally wrong. We are therefore assuming, rather than arguing for, a pre-existing notion of something being morally wrong.¹⁶ We are not questioning whether anything can be morally wrong. Our principle will incorporate this:

EP1 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything on insufficient evidence.

¹⁴ See p74ff.

¹⁵ See also **3.1 Involuntary over-beliefs** p74.

¹⁶ This is in any case implicit in our presupposition of (1.2.1).

Curly brackets { } imply the clause is a contextual presupposition rather than an intrinsic component. Consider the same presupposition in relation to the principle that lying is wrong. A perfectly amoral universe with no concept of moral right or wrong would be one where 'it is wrong to lie' would not apply, because the conditions for its meaningful usage would not obtain.

This does not make every over-belief as wrong as every lie or deception. After all every lie or deception is not equally wrong. But it means that assumptions and distinctions applicable to moral principles in general would, other things being equal, apply to EP1. CP speaks in universal terms, but it would be hard to argue it would be wrong 'always, everywhere, and for anyone' to do *anything* – lie, cheat, steal, or even kill. Morality is a domain of choices and moral imperatives. Taken literally CP says over-believing is the paramount sin: rather sacrifice your only son as a burnt offering than over-believe. But an evidence principle expressed in moral terms need only claim it is wrong to over-believe (as to lie, cheat or steal) in the absence of any conflicting and overriding moral imperative. Only a Kantian extremist would think it morally wrong to lie to the Gestapo about where a fellow maquisard was hiding.

What is meant by 'belief'?

We have distinguished descriptive from prescriptive beliefs,¹⁷ but we also need a working definition of 'belief' itself. For now we will adopt Adler's¹⁸ formula that our beliefs are what we regard as true: 'what I believe is just how things are' (2002:11) for me. This aligns with Williams's claim that beliefs 'aim at truth': '...when somebody believes something, then he believes something which can be assessed as true or false' (1973:136-7).¹⁹

¹⁷ See p14.

¹⁸ See p24ff.

¹⁹ We are therefore taking 'I believe that *p*' as a straightforward way of expressing *p*, and discounting usages where 'I believe that *p*' expresses mere faith, opinion, or a degree of uncertainty that *p*.

An everyday example

We can envisage ostensibly immoral cases of over-belief: dismissing entire races as subhuman because of physical or cultural characteristics, or assuming someone is a terrorist from their appearance.²⁰ But in view of EP1's implicit universality we will pick a more banal example which in itself is not obviously an over-belief:

(1.3) I believe you are taking recreational drugs.

This could be understood as an assertion or as a true self-ascription. As an assertion it would be equivalent to asserting

(1.3.1) You are taking recreational drugs

rather than expressing a degree of uncertainty, which 'I believe' can suggest.

In at least some everyday contexts asserting (1.3) or (1.3.1) without supporting evidence could make me guilty of something like deception. An audience could be justified in assuming I had evidence. If instead I had said:

(1.3.2) I believe you are taking recreational drugs, but I have no supporting evidence,

I would not be deceiving, but might still be blamed for making an unfounded statement or being deliberately mystifying. If (1.3) and (1.3.1) are equivalent, then asserting (1.3.2) is equivalent to asserting

(1.3.3) You are taking recreational drugs, but I have no supporting evidence.

Someone hearing this might struggle to understand what (1.3.2) or (1.3.3) are communicating, since (1.3) and (1.3.1) could be taken to imply the speaker did have evidence. Now consider:

²⁰ See p295.

(1.3.4) I believe you are taking recreational drugs, and I have supporting evidence.

(1.3.5) You are taking recreational drugs, and I have evidence that you are.

These might be heard as making explicit what (1.3) and (1.3.1) leave implicit, and perhaps implying willingness to reveal that evidence.

We are talking here about possible expectations on hearing beliefs expressed. This suggests obligations implicit in belief language itself, possibly supporting why a moral principle like EP1 could seem plausible in respect of our ordinary expectations on hearing everyday beliefs expressed. To help unpack these we will consider in some detail Adler's (2002) 'intrinsic' approach to the ethics of belief, which would ostensibly be opposed to a moral principle like EP1. Our intention is not to incorporate Adler's entire approach into our account, only to co-opt aspects of it to support EP1's implicit universality.

We then move to Haack's (1997, 2001) survey of potential relationships between epistemic and moral appraisals of belief. We suggest an alternative: to see an epistemic evidentialism like Adler's conceptual approach in the social context of belief behaviour, where in turn the conceptual doxastic norms which Adler identifies provide the context and conditions for the kind of moral normativity we see in Clifford.

1.3 Clarifying the principle

Jonathan Adler: *Belief's Own Ethics*

For Adler questions like 'What ought one to believe?' and 'What cannot one believe?' should be answered conceptually (2002:xiii) rather than from a rational,²¹ prudential or – as in EP1 – moral perspective.

²¹ See for example **5.4 Bayesian approaches** p213ff.

His 'intrinsic' or 'basic' ethics of belief claims instead to be 'imposed by the concept of belief itself' (2002:2). He rejects the view that evidence 'should generally determine the strength of belief, but not always': excluding for example 'exciting' beliefs like

supernatural religious beliefs[;] ... beliefs in the goodness or trustworthiness of others; beliefs in fundamental axioms or principles such as induction; [and] beliefs too basic to be supported by anything more certain or fundamental such as that there are external objects. (2002:4)

His main objection to this 'moderate evidentialism' is that, if problematic and unproblematic beliefs have different evidential standards, the difference must be 'determined by factors external to belief': for example, and in particular, rationality. But in respect of unproblematic beliefs the 'demand for adequate reasons or evidence' seems to come from belief itself, not from anything external to belief. If our beliefs are what we regard as true, then 'when we attend to any of our beliefs, a claim is made on us for holding that the belief is true'. So belief itself is what demands 'proportional reasons or evidence', not any external requirement, for example to be rational in our beliefs. Adler does not see why problematic beliefs should make different demands just because their content is different.

We could question though to what extent rationality is 'external to belief'. Davidson for example sees rational (versus non-rational) animals possessing 'propositional attitudes such as belief, desire, intention and shame', which 'come only as a matched set' (1982:317-8).

The issue for Adler would be whether his 'conceptual structure of belief' (2002:2) demanding 'proportional reasons or evidence' for itself (2002:5) can be understood without presupposing rationality: 'it is part of the concept of a belief ... that it tends to cause, and so explain, actions of certain sorts' (Davidson, 2001b:217). Whether or not the conceptual structure of belief presupposes rationality, or any aspect of rationality, is not crucial for us however as EP1 is in any case envisaged as a *moral* principle.

More immediately relevant to EP1 is Adler's preferred 'first-person methodology' (2002:8), focusing on belief from a first-, rather than second- or third-person, perspective:

From the first-person point of view, what I believe is just how things are [for me], not how I conceptualize, interpret, or theorize my experience (2002:11).

Adler's illustration is Jim, who believes he is handsome. Someone else could explain Jim's belief by saying Jim would be depressed if he did not believe it.²² But this non-epistemic reason does not justify Jim's belief the way an epistemic reason would. From his first-person perspective Jim could not acknowledge such a 'disparity between what explains [his] believing and what justifies [his] belief' (2002:5):²³

It would be contradictory for Jim to think that he believes that he is handsome; that his only reason to believe it is that it will lessen his depression; and that the lessened depression does not bear on the truth of whether he is handsome. (2002:9)²⁴

Adler contrasts this with the 'dominant methodology'.²⁵ This 'asks what it is rational to believe' (2002:9), which can only be answered from a third-person perspective:

If Jim's not believing that he is handsome will depress him, without compensating benefit, then it is rational for him to believe it. Yet, from a first-person point of view, Jim cannot take himself to believe that he is handsome as a way to avoid depression. This "cannot" is conceptual.

²² See also p247ff.

²³ There are parallels with James on 'questions concerning personal relations' (1896, 2000:213) (see p110ff), although Jim's 'positive thinking' involves an epistemically unjustified belief about himself rather than an epistemically unjustified belief about another.

²⁴ Adler's example is a little awkward in that Jim's belief that he is handsome could be seen as purely subjective, as could the beliefs of any of Jim's other 'beholders'. We might then dispute if there is any 'truth of whether he is handsome'. To remove this difficulty we will take 'handsome' as equivalent to 'generally considered handsome' – which could then be objectively determined as true or false.

²⁵ Bayesianism? Adler is not explicit. See p213ff.

Jim's (non-epistemic) desire to avoid depression is unrelated to the truth or falsity of the claim that he is handsome. So it cannot be 'the condition of the truth' of his belief that he is handsome. Although it may be rational (as viewed from a third-person perspective) for Jim to believe he is handsome to avoid depression, this in itself would not make it rational for Jim himself, from his own first-person perspective, to see avoiding depression as a reason for thinking his belief is true.

A key aspect of Adler's first-person approach is his 'condition of full awareness', achieved when we '*abstract* from conditions that obscure the concept of belief', and 'bracket the normal unconscious workings of belief and the myriad influences on it'. This full awareness is not an unrealisable ideal but something common to everyday life:

I openly believe both that I believe that it is raining outside and that I do so because I see the rain through the window.

In other circumstances though we may find

[t]he normal condition of believing is that of nonconscious influences and distraction. But when we want to discern what belief demands we should look at it without these interferences, which is accomplished by imposing the full awareness condition. (2002:34)

This condition exposes a 'far-reaching parallel between belief and assertion' (2002:13):

the requirement of the speech act of assertion is to state what is true, as it is the constitutive claim of belief that its content is true.

Adler applies his approach to the kind of belief which anti-evidentialists see as unproblematic, for example: 'I believe that God exists, but by my own lights, I am

not justified in believing that God exists'.²⁶ To assess a claim like this, about a putative belief being unproblematic – if not legitimate – without requiring justification, Adler proposes a four-step 'filtering' (2002:12) process:

Step one: Reserve questioning for the 'privacy of thought' (2002:10), rather than public debate.

Step two: Substitute banal content like 'Tony is in the ice cream parlor' for (eg) 'God exists'.

Step three: Ensure the example is a 'straightforward belief' – simply what I regard as true 'without qualification' – where 'I believe that p ' is equivalent to expressing p , rather than expressing an opinion or partial uncertainty that p (and therefore acknowledging a degree of doubt), or expressing faith or trust in p , both of which can have connotations of personal and/or moral integrity which we aim to 'filter away' (2002:10-11).

Step four: Ask of the belief whether it is really coherent to believe – or assert – that, for example:

(1.4) 'Tony is in the ice cream parlor, but I lack sufficient evidence that it is true' (2002:12).

Adler calls this 'an instance of *Moore's Paradox*'. Both 'Tony is in the ice cream parlor' and 'I lack sufficient evidence that Tony is in the ice cream parlor' could be true. But for Adler asserting (1.4) is as incoherent as asserting ' p , but I do not believe that p ': the paradigm case of Moore's Paradox.

Elsewhere Adler describes conjunctions like (1.4) less controversially as 'Moore's Paradox-like' (2002:30). For example the ancient challenge to believe that

(1.5) 'The number of stars is even' (Burnyeat, 1983:132)

²⁶ Adler (2002:10) quoting Kornblith (1986:119).

cannot be met because this would be incoherent:

(1.5.1) 'The number of stars is even, but I lack sufficient evidence that the number of stars is even' (Adler, 2002:30).

He expands:

(1.5.2) 'I believe that the number of stars is even. All that can secure for me the belief's claim of truth is adequate evidence (reason) of its truth. I lack adequate evidence. So I am not in a position to judge that the number of stars is even. So I do not judge it true. So I do not believe that the number of stars is even.' (2002:30)

We will now apply Adler's four steps to a prosecution lawyer asserting or believing:

(1.4.1) He did it. I know he did it. I just cannot prove it yet.

One: perhaps the lawyer wants to convince his colleagues or impress on himself or others his determination not to rest until he wins the case; but in the 'privacy of thought' his 'I know he did it' could be closer to 'I am *sure* he did it'. Two: could the lawyer's professional reputation and/or the enormity of the crime disqualify the belief from being 'unexciting', 'dull' or 'hardly disputable'? So we should substitute (say): 'His birthday is 6 May (but I cannot prove it)'. Three: is the lawyer taking the content of 'He did it' (or 'His birthday is 6 May') as true 'without qualification'?²⁷ Step four would then ask if it is 'really coherent to believe' that

(1.4.2) His birthday is 6 May, but I lack sufficient evidence that his birthday is 6 May.

A critic could object that Adler's method effectively rules out potential counter-

²⁷ If a detective rather than a prosecution lawyer asserted (1.4.1) we might suspect unprofessional confirmation bias.

examples like (1.4.1) by, crudely speaking, replacing difficult cases with easy ones. More generously however we could see the four steps as a useful way to reveal the core conceptual content which Adler claims is common to all belief. The first-person perspective is just ‘our point of view, ...what I believe is just how things are, not how I conceptualize, interpret, or theorize my experience’:

In belief’s everyday roles, prominently as guides to action, one sees through one’s attitude to the world without seeing that attitude. (2002:11)

Similarly we should abstract assertion from ‘conversational expectations’ while ‘understand[ing] it in its natural home of conversation’. This lets us ‘focus ... on the heart of assertion as the presentation of a proposition as true’ – which ‘exactly parallels the import of belief’ (2002:14).

‘Belief is not up to me’ (2002:14): Adler therefore opposes voluntarism. He agrees with Williams that we cannot see ourselves as holding beliefs because we choose to:

I could not... in full consciousness, regard [something] as a belief of mine, i.e. something I take to be true, and also know that I acquired it at will. (Williams, 1973:148)

This is ruled out by combining ‘one’s beliefs are what one regards as true’ (Adler, 2002:4) with ‘[c]hoice cannot make the belief true’ (2002:15).

But he denies that opposing voluntarism implies we are not responsible for our beliefs (2002:14):

It is undeniable that we can induce a belief by attending to one-sided sources, biasing ourselves toward a sought-for opinion. On the positive side, we can develop dispositions or habits to be open to contrary evidence or criticism, and to evaluate that evidence fairly (2002:56).

EP1, like CP, treats doxastic obligation as a moral obligation, so might

presuppose rather than oppose voluntarism. Adler's ethics of belief is instead 'imposed by the concept of belief itself' (2002:2), not by morality (or, subject to the Davidsonian caveat discussed earlier,²⁸ rationality). Can these be reconciled? Do they need to be? It is time we considered possible relationships between epistemic and moral appraisal of belief.

Susan Haack: *"The Ethics of Belief" Reconsidered*

Haack (1997, 2001) offers several 'serious options' for how epistemic appraisal of belief (as in Adler's intrinsic evidentialism) and ethical (moral) appraisal of belief could relate, and discusses three in detail.

She thinks the 'special-case thesis'²⁹ that 'epistemic appraisal is a subspecies of ethical appraisal' is false (1997, 2001:21-3), because scenarios are possible where a person is epistemically unjustified but not morally at fault. He can only be morally at fault in believing that *p* if, like Clifford's ship-owner, his belief is 'willfully self-induced' (1997, 2001:26). But even if his belief is not willfully induced his evidence may still not be good enough.

EP1 is not necessarily a version of Haack's 'special-case thesis'. Even if this thesis is false – because a person can believe something on insufficient evidence involuntarily and be morally blameless³⁰ – that need not undermine a moral principle that, other things being equal, one should only believe on sufficient evidence.³¹ Aspects of involuntary over-belief will be discussed in **Chapter 3**,³² leading us to adjust EP1.

Next is the 'correlation thesis' (1997, 2001:21) – that positive or negative epistemic appraisal 'is distinct from but invariably' and 'contingently' associated with positive or negative moral appraisal. She offers three possible explanations

²⁸ p25.

²⁹ Associated with Roderick Chisholm, eg (1991).

³⁰ For example person A believes that *p* (but without deliberately choosing to believe that *p*) on the testimony of person B, where A has good reason to trust B's testimony and no reason to distrust it; but it later transpires that B's testimony was false.

³¹ By analogy, someone telling an untruth out of ignorance is not necessarily flouting the moral principle that one should not lie.

³² p74ff.

for unjustified belief: 'negligent incontinence'; 'self-deception'; and 'cognitive inadequacy', either 'personal' or 'cultural'. She sees either kind of cognitive inadequacy as fatal to the correlation thesis even in its weaker 'prima facie' form ('whenever a person believes unjustifiedly, his so believing is always also subject to unfavorable moral appraisal prima facie' (1997, 2001:24)):

If a person has done the best he can, not only to find out whether p, but also to determine that he is competent to find out whether p, he is not morally culpable even if his belief in his competence and his belief that p are, by reason of cognitive inadequacy, unjustified. (1997, 2001:25)

Haack herself favours the 'overlap thesis' (1997, 2001:21): that believing without justification can be, but is not always, a form of 'morally culpable ignorance'. Ignorance itself could be agnosticism; mis-belief ('the belief one has is false'); or over-belief. Of these it is over-belief which 'constitutes culpable ignorance when, as it sometimes but not invariably is, it is both harmful and peccable' (1997, 2001:25-6).

She claims Clifford does not distinguish epistemic from moral culpability, offers no arguments for identifying the two or for the special-case thesis, and just extrapolates his ship-owner story to a universal principle. Clifford's arguments 'could, at most, establish the correlation thesis' (1997, 2001:26).

She lists reasons why the ship-owner's ignorance is morally significant:

The unjustified belief is false; the proposition concerned is of great practical importance; the person concerned is in a position of special responsibility; the false belief leads to dramatically harmful consequences; and the belief is willfully self-induced. (1997, 2001:26)

For the correlation thesis to hold more broadly the ignorance would need to be 'morally culpable even if all these features were absent', say in cases of 'apparently harmless unjustified belief'. Clifford claims that to count as a belief

there must be potential 'influence upon the actions of him who holds it'³³ (1877, 1879:181) – which might then prove harmful. He claims over-belief involves 'doing this great wrong towards Man, that I make myself credulous' (1877, 1879:181) – which for Haack 'carries, if not invariably a risk of harm, a risk of risk of harm' (1997, 2001:27). But she thinks this too 'remote' for moral blame. Otherwise 'not only drunken driving, but owning a car, would be morally culpable'. Also we are not always responsible for unjustified believing: sometimes the cause is cognitive inadequacy (1997, 2001:27).³⁴

The car-ownership analogy seems to miss Clifford's point. His tirade against over-belief covers all consequences: actual and potential; immediate and remote; direct and indirect. Particularly significant is that, like a 'pestilence' or 'plague' (1877, 1879:184), over-belief generates more of itself. The immediate effect of an individual unit of unjustified belief may be harmless, but this is no guarantee that every future and/or indirect effect of every unit of unjustified belief which that original unit may have encouraged (at first, second, third etc remove) will be equally harmless. Car ownership may increase the risk of drunken driving in a statistical sense, in that increasing car ownership might lead to an increase in drunken driving, and dropping car ownership to zero might eliminate it completely. Access to a car is doubtless a necessary precondition of drunken driving. But in itself it does not encourage and/or generate drunken driving. *Drinking and driving* – short of *drunken driving* – is a better analogy though, supporting Clifford's point. Fifty years ago drink driving was more prevalent in the UK, and prevalence normalised it. But 'through firm laws, highly visible enforcement, and a sea-change in public attitudes ... drink driving is now frowned upon by the vast majority of people'.³⁵

The strength of Haack's second objection depends on whether CP implies both additions in square brackets below:

³³ Compare Peirce (1878, 1998:144): 'The essence of belief is the establishment of a habit' – where 'habit' is short for 'rule of action'.

³⁴ We will attempt to immunise our own evidence principle from this kind of objection: see p75.

³⁵ <https://www.gov.uk/government/news/92-of-people-feel-ashamed-to-drink-and-drive-as-50th-anniversary-think-campaign-is-launched> [retrieved 19 December 2017].

CP1.1 *It is [morally] wrong, always, everywhere, and for anyone, to believe anything upon insufficient evidence, [even when the unjustified believing is outside the believer's voluntary control].*

As previously argued,³⁶ Clifford's language strongly suggests he did mean 'morally wrong'. So it is plausible that he did not intend the second addition, and would more likely have accepted its converse:

CP1.2 ... [**except** when the unjustified believing is outside the believer's voluntary control].

This is because once the moral domain is explicit, we can typically assume involuntary acts are excluded. The following seem to be cases of linguistic incomprehension:

It is morally wrong to lie, except when you had every reason to think you were telling the truth or when talking in your sleep.

It is morally wrong to deceive someone, except when the person had misinterpreted an act [statement] you had performed [made] in innocence.

'Lying' typically excludes innocent untruths, and 'deceiving' typically excludes being misinterpreted. Haack agrees 'bad habits' of unjustified belief 'may, if unchecked, become inveterate', yet doubts if 'indulgence in such habits is bound to encourage them in others' (1997, 2001:27). But Clifford need not claim every over-belief is 'bound to' inch us into savagery, which is why his 'plague theory' (Gale, 1993:356) seems so apt. An HIV+ individual refusing to practise safe sex is not 'bound to' infect every partner, but this does not exonerate him.

Haack's list of ways epistemic and moral appraisal of belief could relate includes hierarchical (special case) and empirically discovered relationships (correlation and overlap); 'independence' (no relationship); and 'analogy'. But even her

³⁶ pp14–17.

favoured overlap thesis incorporates no account of how and why they may overlap. We now consider another option: that a conceptual (therefore epistemic) normativity like Adler's may set the context and conditions for the kind of moral normativity Clifford identifies.

Social context of belief: 1

Adler sees violations of doxastic norms (like 'p, but I lack adequate evidence that p') as conceptually 'incoherent' (2002:3) rather than immoral or irrational.³⁷

Clifford stresses however that acquiring, holding, asserting and acting on beliefs are straightforwardly and inextricably social. Belief helps 'bind men together' and 'strengthen and direct their common action' (1877, 1879:182-3).³⁸ **Chapter 3**³⁹ will ask whether truly private and unvoiced beliefs are problematic for an evidence principle. But such beliefs aside (and for argument's sake also restricting ourselves to beliefs surviving Adler's 'filtering' process⁴⁰) if doxastic norms legislate which conceptual moves may or may not be made when we believe, such norms will contribute to the context of expectations where socially contextualised belief can and will be morally assessed.

As an analogy, shared understanding of arithmetic contributes to the context of expectations where examination results would be calculated, communicated and acted on. A student getting 25% might speak and behave as if he was 'top of the class' – because in his mind 25 was higher than 90. He might invite accusations of self-deception and/or stubborn denial and/or disrespect – particularly from colleagues who had laboured all year to improve their marks.

The student 'ought' to see 25% is lower than another's 90%. At the equivalent of Adler's 'basic' (2002:3) level this is not a moral ought but a conceptual ought derived from the shared rules of arithmetic. But a moral ought can arise when evaluating social behaviour within the context of that conceptual normativity. We

³⁷ Subject to Davidsonian caveat: p25.

³⁸ See also p210ff.

³⁹ See p122ff.

⁴⁰ p28.

would exonerate someone whose learning disability stopped him applying the rules of arithmetic. But we might blame someone who seems capable of understanding the rules but now apparently wants to violate them to bask in undeserved glory.

Or consider truth and lying. Whether a statement is true or false, and what makes it true or false, are matters for (say) science or logic, not morality. But a statement's capacity to be true or false makes lying and telling the truth possible, and these have moral significance. The conceptual evidentialism which Adler derives by analysing what having and expressing a belief involve is an epistemological, not a moral, thesis. But assuming Adler is correct, the interpersonal expectations grounded in shared assumptions about how beliefs, assertions and evidence interrelate conceptually can expose someone asserting 'I believe you are taking recreational drugs' (1.3) or 'You are taking recreational drugs' (1.3.1) without supporting evidence to a charge of deception.⁴¹

Every descriptive over-belief will not have equally severe moral repercussions. But holding and asserting a descriptive over-belief *can* have moral repercussions – directly related to the lack or inadequacy of supporting evidence, against a background expectation that holding and/or asserting a belief typically implies the believer does have sufficient evidence.

Below is an Adlerian analysis of (1.3.1):⁴²

If I believe and/or assert (1.3.1), then I regard (1.3.1) as true. When I attend to belief (1.3.1) a claim would be made on me for holding (1.3.1) as true (Adler, 2002:4). The link between the claim of the truth of (1.3.1) and the condition of the truth of (1.3.1) is whatever evidence I have for (1.3.1). If I fully believe that (1.3.1) I need adequate epistemic reasons – evidence – for thinking (1.3.1) is true.

⁴¹ This expands the point summarized in (1.1.6): p19 above.

⁴² Following Adler we will take (1.3) ('I believe you are taking recreational drugs') and (1.3.1) ('You are taking recreational drugs') as equivalent.

As a parallel to Jim's belief that he is handsome, someone else could explain my belief (1.3.1) that you are taking drugs by (say) my desire to shine at your expense by casting you as reckless and immature. But while this non-epistemic reason might explain my belief it cannot justify it. I cannot in 'full awareness'⁴³ take myself to be holding belief (1.3.1) as a way of shining at your expense – as Jim cannot take himself to believe he is handsome so as to avoid depression (2002:10).

This 'cannot' is primarily a conceptual, not a practical or moral, constraint. Adler would see it as 'contradictory' for me to think that I believe you are taking drugs; *and* that my only reason for believing this is to shine at your expense; *and* that my elevated self-esteem has no bearing on the truth of whether you are taking drugs (2002:9). My intention or desire to shine cannot be 'the condition of the truth' (2002:5) of my belief that you are taking drugs.

We could question whether it is exactly contradictory, but it is at least conceptually unstable. Reflecting on the belief and what could be the condition of its truth should lead the believer to discard the belief.

Adler would say that asserting '(1.3.1), but I lack adequate evidence that (1.3.1)⁴⁴ violates a conceptual doxastic norm and is therefore conceptually incoherent (2002:3). But the shared doxastic norms legislating which conceptual moves may or may not be made are part of the shared context of expectations within which social interaction – including assertion and reception of beliefs like (1.3.1), and behaviours resulting from those beliefs – might be morally assessed. It is that shared context of expectations and assumptions which makes asserting (1.3.1) without sufficient evidence equivalent (or close or akin) to deception, and therefore morally culpable.

But even if conceptual normativity does provide a context for moral normativity, why is conceptual normativity not enough for an ethics of belief? Why do we need

⁴³ See p38ff.

⁴⁴ Equivalent to:
(1.3.3) You are taking recreational drugs, but I have no supporting evidence.

moral normativity?

Why both?

Take Adler's 'full awareness' – or Williams's equivalent 'full consciousness' (1973:148). Saying this condition needs to be imposed or assumed implies a person could acquire or hold a belief without it, which in turn suggests (moral) choice:

*There is an **ethical** grounding to the condition of full awareness that supports this first-person methodology. If one is acting (believing) rightly, one can acknowledge it without **guilt or shame**. [Emphases added.] (Adler, 2002:9).*

The person apprehending, describing, conceiving or evaluating the situation is the subject whose situation it is. In the context of (1.5.1) and (1.5.2) the person who is, or can be, aware of having insufficient evidence that the number of stars is even is the subject believing or asserting that the number of stars is even. If he also satisfies the condition of full awareness then he is taking personal responsibility for the implications of believing the number of stars is even and of knowing no evidence is available. The subject taking responsibility is then faced with the incoherence. A subject avoiding the condition of full awareness by not taking responsibility is not faced with the incoherence. Being aware or unaware of the incoherence would indicate whether or not one was meeting the condition of full awareness. The first-person perspective is not a matter of choice, but applying full awareness can be. Adler's 'ethical grounding' equates to an implicit moral duty to be fully aware, or as aware as reasonably possible.

From a third-person perspective there is no incoherence:

(1.5.3) Fred believes the number of stars is even.

(1.5.4) Fred has no evidence the number of stars is even.

Both statements could be true – although from a third-person perspective someone can appreciate the incoherence Fred is faced with.

The first-person perspective and the condition of full awareness are also relevant to Adler's third step: ensuring it is a 'straightforward belief' that p , rather than expressing faith in p , or an opinion or partial uncertainty that p (2002:10). A test of full awareness would be whether the subject, from the first-person perspective (with its implication of taking or not taking personal responsibility), can tell if his belief is a 'straightforward belief', in respect of which, as with an assertion, 'a claim is made on us for holding that the belief is true' (2002:4). This seems to be where the implicit moral duty comes in.

We should avoid a potentially misleading implication though. A third-person perspective necessarily entails a social context since (for example) someone other than Jim is explaining Jim's belief that he is handsome. A first-person perspective on the other hand might theoretically involve just the believer – Jim, say – with no implied social context. But believing is social behaviour (Clifford, 1877, 1879:182-3) regardless of perspective. From what Adler calls the third-person perspective the focus is on interpretation and explanation: a third party speculates (say) about why Jim believes as he does, and suggests reasons which may be epistemic or non-epistemic. Any obligations Jim might acknowledge, or fail to acknowledge, have not necessarily entered the picture, unless the third party is also (morally or otherwise) evaluating the believer's behaviour.

The first-person perspective on the other hand necessarily involves the believer's obligations. Under Adler's condition of full awareness these obligations arise purely from the 'conceptual structure of belief' (2002:2). But the minute we acknowledge the social nexus within which doxastic behaviour operates, those conceptual obligations flesh out into moral obligations. This is why conceptual normativity is not enough for an ethics of belief: because believing is intrinsically social behaviour, the conceptual implications of what we assert or reveal about our beliefs will generate expectations which we can either honour or shirk.

So our conclusion is that even if Adler's 'intrinsic' ethics of belief applied

comprehensively and soundly to all categories of belief, it does not remove the rationale for a moral evaluation of belief in its social context. Clifford seems therefore to have got something fundamentally right. We may agree with Haack that Clifford does not clearly distinguish epistemic from moral culpability. But he does not conflate the two. He sees doxastic failings as moral failings, at least where a degree of volition is involved. But he does not offer a contextual analysis like the one just proposed, aimed at separating those aspects of belief with moral significance from those without.

Epistemic obligation may not be a strict subset of moral obligation, as in Haack's 'special-case thesis'. But epistemic and moral obligation do not just 'correlate' or 'overlap' either. Some aspects of epistemic obligation appear decidedly moral and others might be more appropriately considered in purely conceptual terms. But in a social context they are inextricably linked, such that spelling out doxastic obligation as a guiding principle can require moral language. Adler's approach does not remove the rationale for an evidence principle expressed in moral terms.

This might therefore be a convenient place to address a potential objection that a moral principle like EP1 is by its very nature irrelevant to epistemology. We have said EP1 is primarily intended as a self-standing moral principle applying to an agent's doxastic behaviour, and only secondarily to the resultant beliefs themselves.⁴⁵ It is not furthermore separately intended as a comprehensive criterion of epistemic justification derived from moral considerations, if this implies a moral normativity which is ultimately a means to an end, that end being epistemic normativity. We would question though why this should make EP1 irrelevant to epistemology, except on a deliberately and unnecessarily narrow definition of 'epistemology'. Belief, believing, and believing language do not occur in a vacuum but within a social, and therefore potentially moral, context. This is what both CP and EP1 recognise.

⁴⁵ See p21.

1.4 Reasons for holding an evidence principle

An evidence principle expressed in moral terms lays on the believer a prima facie burden of justification for acquiring and/or holding descriptive beliefs. To justify a belief morally, the believer would need to demonstrate either that it is supported by evidence, or that by not believing he would breach a conflicting and overriding moral imperative.⁴⁶ Without evidence, the content of the descriptive belief cannot be used to justify it. This has particular relevance for religious belief.⁴⁷

An evidence principle expressed in moral terms will inherit assumptions and distinctions applicable within ethics generally.⁴⁸ It would therefore need to 'fight its corner' since an actual instance of over-belief could be exonerated by a conflicting and overriding moral imperative. The same applies to principles against cheating, lying, stealing – even killing.

It might however be over-optimistic to expect a completely comprehensive justification for an evidence principle expressed in moral terms. Earlier⁴⁹ we articulated a rationale for CP which EP1 would largely share, including a set of empirical claims (1.1.1–13) plus a generic moral claim (1.2.1) which, although ostensibly plausible, we could only presuppose rather than establish independently. With many if not all moral principles it seems possible to defend a counter-position – particularly if 'thesis' and 'antithesis' come from opposed meta-ethical perspectives.⁵⁰

So perhaps the most we can hope for is that our principle makes sense as a moral principle; that it behaves like one and entails no contradictory or intuitively unacceptable consequences; and that it can be positioned in relation to the kinds of meta-ethical position it might be compatible or incompatible with.

⁴⁶ See p22.

⁴⁷ See p137ff.

⁴⁸ See p22ff.

⁴⁹ pp18–21.

⁵⁰ For example Thrasymachus's claim in Plato's *Republic* that 'justice or right is simply what is in the interest of the stronger party' contradicts an intuitive notion of justice as fairness. (Plato, 1955, 1974 (Revised)§338c p77).

1.5 Conclusions

This introductory chapter illustrated how Clifford presents and promotes CP: as a moral principle, therefore prioritising voluntary rather than involuntary aspects of belief acquisition. He applies CP to both descriptive and prescriptive beliefs and his defence focuses on the direct and indirect social impact of obeying or disobeying it.

We then made a prima facie case for a moral principle EP1 based on Clifford's but restricted to descriptive beliefs. Like Clifford we focused on social context, but unlike Clifford we looked at how belief and evidence relate conceptually, by considering what our everyday expectations might be when using and encountering belief language. Adler's (2002) 'basic' ethics of belief ostensibly eschews rational or moral grounds for evidentialism, so in theory might oppose a Clifford-style moral principle. Moving to Haack's (1997, 2001) survey of how epistemic appraisal relates to moral appraisal, we broadly endorsed her 'overlap' thesis.

Haack says little however to explain the overlap. We argued that even if Adler's case is plausible, his 'methodology' involves an implicit moral stance; and when the social context of belief is acknowledged his conceptual epistemic normativity can supply the conditions for moral normativity. More importantly, a conceptual evidentialism like Adler's renders a moral principle like EP1 neither invalid nor irrelevant.

We ended by briefly positioning potential reasons for holding a moral principle requiring evidential support for descriptive beliefs. The next chapter argues why our principle should exclude prescriptive beliefs.

Chapter 2 | Prescriptive beliefs

2.1 Introduction

This chapter explores whether an evidence principle like CP can apply to prescriptive beliefs. It also introduces Clifford's perhaps most celebrated opponent, William James. The distinction between descriptive and prescriptive beliefs is relevant to their debate for a number of reasons.

First, both Clifford¹ and James use examples of both. James's initial examples come in a possibly ironic list:

Here in this room, we all of us believe in molecules and the conservation of energy, in democracy and necessary progress, in Protestant Christianity and the duty of fighting for 'the doctrine of the immortal Monroe'. (1896, 2000:203)

Belief in the 'duty' of fighting for the Monroe Doctrine is unambiguously prescriptive, while beliefs in 'democracy', 'necessary progress' and 'Protestant Christianity' could be either.

Second, CP is itself prescriptive – indeed morally prescriptive in Clifford's presentation. Even those who think CP more plausible as an epistemic than a moral principle would see holding it as holding a second-order prescriptive belief covering both prescriptive and descriptive first-order beliefs.

Third, James doubts CP's application to moral beliefs (1896, 2000:212ff). And finally religious beliefs, certainly as James defends them, typically have prescriptive components.² This would be another reason to explore whether CP may apply less successfully to prescriptive than to descriptive beliefs.

We have two potential conclusions in sight. The stronger is that a principle like CP

¹ See p14 and also p52ff.

² See p131ff.

cannot apply to prescriptive beliefs. The weaker is that excluding prescriptive beliefs from a principle like CP is justifiable and not arbitrary, because significant issues arise when applying it to prescriptive beliefs which do not arise with descriptive beliefs. The weaker conclusion is sufficient for our purposes, because our aim is to see what Clifford has got fundamentally right, even if this means modifying his original scope.

Chapter 1 defined a prescriptive belief as being about what ought (or ought not) to be the case.³ We should now be more specific.

2.2 Hypothetical prescriptive beliefs

We must distinguish categorical from hypothetical prescriptive beliefs. Hypothetical prescriptive beliefs should not be problematic for an evidence principle because they can generally be reworded as descriptive beliefs without significant loss of meaning.

Clifford mentions for example the 'Australian' who insists on tying his hatchet blade to the side of the handle, despite the hole provided, because he has 'sunk so low' he cannot 'call in question an established usage, and invent or learn something better' (1877, 1879:203). The context is therefore prudential beliefs about right and wrong ways to do things:

(2.1.1) You should fit the handle into the hole in the hatchet rather than tie the two together.

Evidence supporting (2.1.1) is the success achieved when the handle is fitted into the hole rather than when the hatchet is tied to the side of the handle. If the evidence were different the belief would not be justified.

But the evidence here supports underlying descriptive beliefs: that a hatchet can be struck harder and more accurately when the handle is fitted in the hole and

³ See p14.

gets detached from its handle if it is just tied. The apparently prescriptive belief (2.1.1) assumes an intention:

(2.1.2) **[If you want to chop wood with less effort and less risk of injury then]** you should fit the handle into the hole in the hatchet, not just tie the two together.

The evidence for (2.1.2) might be:

(2.1.3) A hatchet can be struck harder and more accurately when the handle is fitted in the hole; and a hatchet gets detached from its handle if it is just tied on.

But (2.1.3) only supports (2.1.1) given the intention in (2.1.2). If the intention was different (2.1.3) would not be supporting evidence:

(2.1.4) **[If you want to carry out the stage directions for the character you are playing then]** you should fit the handle into the hole in the hatchet, not just tie the two together.

The evidence for (2.1.4) might be the stage directions themselves, not (2.1.3).

Hypothetical prescriptive beliefs like (2.1.2) and (2.1.4) can be reworded as equivalent descriptive beliefs:

(2.1.2d) Fitting the handle into the hole in the hatchet rather than tying the two together will allow you to chop wood with less effort and less risk of injury.

(2.1.4d) The stage directions say the character you are playing fits the handle into the hole in the hatchet rather than tying the two together.

(2.1.1) is shorthand for an explicit hypothetical prudential belief like (2.1.2). If (2.1.1) was a truly standalone categorical prescriptive belief with no assumed

intention as in (2.1.2) or (2.1.4) then it is unclear what evidence could support it, as (2.1.2) and (2.1.4) require different evidence.

From now on 'prescriptive belief' will always indicate a categorical prescriptive belief unless otherwise stated.

2.3 Evidence and truth value

Consider (categorical) prescriptive variants of (1.3) and (1.3.1):⁴

(2.2) I believe you should not be taking recreational drugs.

(2.2.1) You should not be taking recreational drugs.

Again the context is where the speaker has no supporting evidence, so could have been more explicit:

(2.2.2) I believe you should not be taking recreational drugs, but I have no supporting evidence.

(2.2.3) You should not be taking recreational drugs, but I have no supporting evidence.

We are taking (2.2) and (2.2.1–3) as purely prescriptive. We are assuming the person addressed is taking recreational drugs, and ignoring any existent or non-existent evidence for related descriptive beliefs, for example about the effects of drug-taking. In such a context it appears less straightforward that we would blame someone for holding a prescriptive over-belief than for holding a descriptive over-belief. This is assuming it is even meaningful to speak of evidence supporting a prescriptive belief.

⁴ See p23.

Reasons for a prescriptive belief

Consider a variant of (2.2) about a specific drug D, where there is no evidence of any medical or psychological ill-effects, therefore sufficient evidence there are none. In spite of that balance of evidence it seems I could still say to someone, without sufficient evidence,

(2.2.D) I believe you should not be taking D

without being guilty of deception or of making an unfounded accusation. I could be considered judgmental, but this comes of voicing the prescriptive belief itself. It is unrelated to whether or not I have sufficient evidence for (2.2.D), assuming there could be any.

I could have reasons for (2.2.D) which do not count as evidence. We have previously distinguished between epistemic and non-epistemic reasons in relation to descriptive beliefs.⁵ A non-epistemic reason for (2.2.D) might be that I dislike seeing other people enjoying themselves. Someone else could believe or assert this about me:

(2.2.DR1-3) He believes you should not be taking D because he dislikes seeing other people enjoying themselves.

In the third person (2.2.DR1-3) would be a legitimate and reasonable statement. But would equivalent propositional content fail Adler's 'first-person' test, as it would if (2.2.D) was a descriptive belief? I myself could believe or assert:

(2.2.DR1-1) I believe you should not be taking D because I dislike seeing other people enjoying themselves.

Applying Adler's formulation (2002:9):

(2.2.DR1-1A1) It would be contradictory for me to think that I believe you should

⁵ See p26ff.

not be taking D; that my only reason to believe it is because I dislike seeing other people enjoying themselves; and that my disliking seeing other people enjoying themselves does not bear on the truth of whether you should not be taking D.

Any 'contradiction' here is less obvious than with a descriptive belief – if only because of the challenge of interpreting 'the truth of whether you should not be taking D'. In Adler's descriptive belief example the gulf is between Jim's subjective desire to avoid depression and the objective truth or falsity⁶ of his handsomeness. But it is not obvious that my judgment (2.2.D) can be objectively true or false. In the case of (2.2.DR1-1) it might be safer to see a gulf between the subjectivity of my disliking seeing other people enjoying themselves and the (perhaps implicit) universality of my moral judgment that you should not be taking D. My 'contradiction' would be failing to see the incongruity of using a personal preference to justify a moral judgment. By comparison the unproblematic third-person version (2.2.DR1-3) is telling a possibly true story about that failure.

For a prescriptive belief like (2.2.D) a safer (because more evasive) formulation might be:

(2.2.DR1-1A2) [...] and that my disliking seeing other people enjoying themselves has no bearing on whether I am right that you should not be taking D.

Now consider another possible reason:

(2.2.DR2-1) I believe you should not be taking D because I think it wrong to nurture inauthentic behaviour [or to be self-indulgent, or to lose self-control]

Applying our reformulated 'Adler test':

⁶ See footnote p26.

(2.2.DR2-1A1) It would **not** be contradictory for me to think that I believe you should not be taking D; that my only reason to believe it is because I think it wrong to nurture inauthentic behaviour [or ...]; and that the wrongness of nurturing inauthentic behaviour [or ...] **does** have a bearing on whether I am right that you should not be taking D.

There is no 'contradiction' here because it is not incongruous to use one prescriptive belief ('it is wrong to nurture inauthentic behaviour') as a moral reason to justify another if they are appropriately related. But it is another question whether a moral reason qualifies as evidence (an epistemic reason) for (2.1.D). So although we have found a kind of reason which passes an amended version of Adler's 'first-person' test we are no closer to understanding whether evidence can support a prescriptive belief.

Truth value of a prescriptive belief

We reworded Adler's test to avoid unpacking 'the truth of whether you should not be taking D'. But if evidence for a descriptive belief is a reason for thinking the belief is true, then we must ask if a (categorical) prescriptive belief is the kind of thing which can be true or false. If a prescriptive belief p cannot have a truth value⁷ then perhaps we can only have reasons for holding p which do not count as evidence because they cannot count as reasons for thinking p is true – because p is not the kind of thing which can be true. Opinions about inauthentic behaviour, self-indulgence or loss of self-control might then be (non-evidential) reasons for holding a prescriptive belief without being reasons for thinking the belief is true. We might for example think our prescriptive beliefs must be consistent with each other.

But this jars with our working definition of 'belief'.⁸ We accepted 'when we attend to any of our beliefs, a claim is made on us for holding that the belief is **true**'

⁷ We are not assuming that if something is a belief then it has a truth value. See discussion below about vague descriptive beliefs.

⁸ p22.

(Adler, 2002:4); and 'beliefs aim at **truth** ... [so] when somebody believes something, then he believes something which can be assessed as **true or false**' (Williams, 1973:136-7). [Emphases added.] Does our working definition only apply to descriptive beliefs?

If Adler's and Williams's statements must apply for something to qualify as a belief, we cannot use 'belief' to refer to the content of just any statement employing the verb 'believe'. Yet our discussion so far presupposes that 'prescriptive belief' refers to something legitimate: that, for example, 'I believe killing is wrong' expresses a prescriptive belief that killing is wrong while 'I believe the moon is made of green cheese' expresses a corresponding descriptive belief. But if a prescriptive belief is not the kind of thing which can have a truth value then our Adler/Williams working definition would disqualify it from being a belief at all – because a belief by definition must have a truth value, whether we know it or not.

It might appear uncontentious that descriptive beliefs have a truth value. But there are possible exceptions, for example vague descriptive beliefs.⁹ With prescriptive beliefs though we can ask if there might be general reasons for thinking no prescriptive belief has a truth value.¹⁰ It seems possible to acknowledge arguments on both sides: ones supporting the claim that, say, (i) the belief that killing is wrong *does* have a truth value; and ones supporting the opposing claim that (ii) the belief that killing is wrong does *not* have a truth value. Claim (i) might get support from a redundancy theory of truth, holding that 'killing is wrong' and 'it is true that killing is wrong' are equivalent. Claim (ii) might be supported by a verification theory of truth, holding that a proposition can only be true or false if it can be verified, and it is unclear whether 'killing is wrong' can be verified.

For now we will continue assuming 'prescriptive belief' has legitimate reference, and note that our working definition may be too restrictive.

But doubts about prescriptive beliefs having a truth value could make us question

⁹ An example might be the belief that Joaquin Phoenix is short. The actor is 1.73 metres. See previous footnote.

¹⁰ Excluding explicit hypothetical prudential beliefs like (2.1.2) and (2.1.4).

whether applying an evidence principle like CP to prescriptive beliefs is coherent. Someone adopting CP would hold it as a prescriptive belief. But if prescriptive beliefs cannot have a truth value then we could not have reasons for holding it to be true. We might have reasons for holding a prescriptive belief – as for example in (2.2.DR2-1) – but not reasons for holding it to be true.¹¹

Yet if:

- (i) to hold CP is to hold a prescriptive belief; and
- (ii) a prescriptive belief cannot have a truth value; and therefore
- (iii) there cannot be a reason for thinking it is true; and therefore
- (iv) there cannot be evidence for it;

then if:

- (v) CP applies to prescriptive beliefs; and therefore
- (vii) CP applies to the belief that CP;

then:

- (viii) to hold CP is to hold a belief we have no right to hold.

We could also ask if it is coherent to hold CP without holding it to be true. But we could ask the same about moral principles against lying or stealing.

Chapter 1 made a prima facie case for moral culpability for holding and/or asserting at least some kinds of descriptive over-belief.¹² We have not found an

¹¹ Someone might of course see herself as having a reason for holding a prescriptive belief to be true because she believes (possibly wrongly) that a prescriptive belief can have a truth value.

¹² See p21ff.

equivalent prima facie culpability in the case of prescriptive beliefs. We could have reasons for holding a prescriptive belief which do not count as evidence. Or, if a prescriptive belief is not something which can have a truth value, and evidence is a reason for holding a belief to be true, then supporting evidence may not even be possible for prescriptive beliefs. Alternatively, even if prescriptive beliefs can have a truth value they still may not be the kind of thing which can have supporting evidence.¹³

Clifford thought differently though. We should take a closer look at how he treated morally prescriptive beliefs in particular.

2.4 Clifford on moral beliefs

Clifford clearly applied CP to moral beliefs, talking of the 'care and honesty' with which 'great principles, ... most fitted for the guidance of life', have been 'tested', thereby acquiring a 'practical certainty':

The beliefs about right and wrong which guide our actions ... never suffer from investigation; they can take care of themselves, without being propped up by "acts of faith", the clamour of paid advocates, or the suppression of contrary evidence. (1877, 1879:188-9)

He sees the 'time-honoured tradition of the human race' (1877, 1879:199) as a 'great fabric'

for the guidance of our thoughts, and through them of our actions... . In the moral world ... it gives us the conceptions of right in general, of justice, of truth, of beneficence These are given as conceptions, not as statements or propositions; they answer to certain definite instincts, which are certainly within us, however they came there. (1877, 1879:201)

¹³ **Appendix A2** (p281ff) argues that if a prescriptive belief is not something which can be supported by evidence, it may not be able to feature in 'Bayesian updating'.

'Conceptions' which are not 'statements or propositions' may also be incapable of being true or false, and therefore not supportable by evidence. But what Clifford seems to see as subject to evidence is not the conceptions themselves but our understanding of how to apply them: for example how 'right' relates to 'beneficence':

That it is right to be beneficent is matter of immediate personal experience; for when a man retires within himself and there finds something, wider and more lasting than his solitary personality, which says, 'I want to do right,' as well as, 'I want to do good to man,' he can verify by direct observation that one instinct is founded upon and agrees fully with the other.

Without using the word he does seem to be referring to beliefs: that wanting to do good to man is founded on wanting to do right and that doing good to man is a way of doing the right thing.

He also shifts from 'is' to 'ought'. If 'direct observation' reveals the connection between the instinct to do right and the instinct to do good to man, that implies 'doing good to man' can be independently characterised in non-normative terms (say by quantifying maximum net happiness to human beings¹⁴) without 'doing good to man' necessarily being the 'right' thing to do. Otherwise there would be nothing to 'observe', as 'doing good to man' would necessarily be the right thing to do. Clifford seems to be saying that as a matter of fact

(2.3.1) people have the instinct to 'do right'; and

(2.3.2) they have the instinct to 'do good to man'; and

(2.3.3) they observe a connection between (2.3.1) 'doing right' and (2.3.2) 'doing good to man' – such that doing good to man is a way of doing the right thing.

¹⁴ Although see **Appendix A1** p276ff for possible limits to Clifford's utilitarianism.

(2.3.1–3) together are the evidence for believing ‘it is right to be beneficent’; or in other words:

(2.4) We ought to be beneficent.

But it is our ‘duty ... to verify this and all similar statements’ (1877, 1879:202). So how do we respond if someone claims not to experience the connection, so does not have the same evidence? Assume my belief that we ought to be beneficent applies to people in general and not just myself. Someone could doubt I really have the evidence I claim for believing people in general should be beneficent. She may accept I have sufficient evidence for believing it *feels* right to be beneficent, or believing it is right *for me* to be beneficent. But if I have only subjective experience of (2.3.1–3) occurring together this seems insufficient evidence for believing it is right *for people in general* to be beneficent (2.4). In Clifford’s own language: how do I know the something I find within myself is ‘wider and more lasting than [my] solitary personality’?¹⁵

A sociopath could claim not to experience any of (2.3.1–3). A religious fundamentalist might experience (2.3.1) but not (2.3.2) or (2.3.3) because for her ‘doing right’ equals doing what her God commands, which may exclude ‘doing good to man’.

But Clifford takes it as established that it is right to be beneficent (and/or just, and/or truthful) because our ‘moral sense founded on experience’ (1877, 1879:202) provides sufficient evidence. He then asks how to know ‘true beneficence’ – how to be as practically beneficent as we can in our ‘definite place and time’. We are again in the domain of beliefs with possible truth values. In respect of beneficence,

the great social heirloom consists of two parts: the instinct of beneficence, which makes a certain side of our nature, when predominant, wish to do good to men; and the intellectual conception of beneficence, which we can

¹⁵ See p226ff, on the need for evidence to be ‘objective’, or at least publicly available. Clifford elsewhere asks a related question about the ‘best conscience’: see **Appendix A1** p276ff.

compare with any proposed course of conduct and ask, "Is this beneficent or not?"

We therefore need a 'further inquiry' to determine how, say in utilitarian¹⁶ terms, to generate maximum net happiness. Clifford gives an example in his claim that re-examining the evidence led to a rethink about the treatment of beggars:

Until recently, the moral tradition... taught that it was beneficent to give money indiscriminately to beggars. But the questioning of this rule, and investigation into it, led men to see that true beneficence is that which helps a man to do the work which he is most fitted for, not that which keeps and encourages him in idleness; and that to neglect this distinction in the present is to prepare pauperism and misery for the future.

This appears to reference a social policy shift reflected in English Poor Law legislation. The 'before' and 'after' beliefs could be summarised as, respectively:

(2.5) We ought to give money indiscriminately to beggars.

(2.6) We ought to help people do the work they are most fitted for, and not keep and encourage them in idleness.

Both are morally prescriptive, and Clifford's treatment identifies the instinct to 'do right' (2.3.1) as the ultimate source of their prescriptive element. As humans we experience ourselves as instinctively moral beings (2.3.1) and as instinctively beneficent beings (2.3.2); and we also experience how the two connect (2.3.3). Even if Clifford is right about this, and therefore right about our sense of obligation (2.3.1), this does not explain how the sense of obligation we happen to experience ('is') generates our actual obligation ('ought'). Perhaps he sees them as identical? An earlier essay of his could be interpreted as a 'social evolutionist' approach to the emergence of morality, suggesting he saw moral evaluation as in principle like

¹⁶ But see p68ff.

finding the best way to build a bridge. If so this might in turn explain his application of CP to prescriptive beliefs. We explore this in **Appendix A1**.¹⁷

Giving money indiscriminately to beggars (2.5) and helping people do the work they are most fitted for (2.6) are both potential ways of being beneficent. So (2.5) and (2.6) both presuppose a more generic (categorical) morally prescriptive belief:

(2.4) We ought to be beneficent.

For convenience we will interpret 'beneficence' using a descriptive utilitarian formula like maximising net happiness. In principle this could help identify the evidence to decide between (2.5) and (2.6) – as long as we agree the formula. But even if we can find utilitarian evidence favouring (2.6) over (2.5) or vice versa, this will not necessarily qualify as the kind of 'evidence in support of a prescriptive belief' we need. This is because if (2.5) and (2.6) presuppose categorical prescriptive belief (2.4), this allows us to interpret beneficence in non-normative terms, therefore to treat the 'social policy' content of (2.5) and (2.6) as purely prudential (so hypothetical) prescriptive beliefs.¹⁸

(2.5h) [If we intend to be as beneficent as possible then] we ought to give money indiscriminately to beggars.

(2.6h) [If we intend to be as beneficent as possible then] we ought to help people do the work they are most fitted for, and not keep and encourage them in idleness.

We could reword these into descriptive beliefs without significant loss of meaning:

(2.5d) Giving money indiscriminately to beggars would maximise beneficence.

¹⁷ See p276ff.

¹⁸ See p44ff.

(2.6d) Helping people do the work they are most fitted for, rather than keeping and encouraging them in idleness, would maximise beneficence.

Any balance of evidence allowing us to choose between (2.5d) and (2.6d) – therefore between (2.5h) and (2.6h); and therefore between (2.5) and (2.6) on the presupposition of categorical prescriptive belief (2.4) – would be evidence supporting descriptive, not prescriptive, beliefs.

Clifford's 'social evolutionist' utilitarian approach¹⁹ may explain his application of CP to prescriptive beliefs – if he saw moral evaluation as in principle like deciding how to build a bridge. But this does not explain how evidence could support a categorical prescriptive belief like (2.4). There may be evidence that the evolution of beneficence promotes the survival of the group which enjoys it but this can only be evidence that the group ought to be beneficent if there is also evidence that the group ought to survive. Nor though have we established that evidence *cannot* support a categorical prescriptive belief.

We will now see what James's more individualist perspective can offer.

2.5 William James on prescriptive beliefs

Moral beliefs are James's first counter-example to CP, because they

cannot wait for sensible proof. A moral question is a question not of what sensibly exists, but of what is good, or would be good if it did exist. Science can tell us what exists; but to compare the worths, both of what exists and of what does not exist, we must consult not science, but what Pascal calls our heart. (1896, 2000:212)

The heart reference reflects a celebrated quotation a few paragraphs back:

¹⁹ See p276ff.

*Human passions, however, are stronger than technical rules.*²⁰ “*Le cœur a ses raisons,*” as Pascal says, “*que la raison ne connaît pas;*”... (1896, 2000:212)

The quotation continues:

423 *The heart has its reasons of which reason knows nothing: we know this in countless ways.*

I say that it is natural for the heart to love the universal being or itself, according to its allegiance, and it hardens itself against either as it chooses. You have rejected one and kept the other. Is it reason that makes you love yourself? (277)

424 *It is the heart which perceives God and not the reason. That is what faith is: God perceived by the heart, not by the reason. (278) (Pascal, 1966:154)*

If the Pascalian heart is something which ‘love[s] the universal being or itself, which ‘hardens itself against either as it chooses’, and ‘perceives God’, then not everyone would recognise that in their lives – in James’s 1890s at least. If, more broadly, it is something which loves either itself or something else which is not the self, then we seem to be talking about orientation: perhaps an aspect of the self which determines whether the self is self-oriented or oriented towards others (or towards the ‘other’; or the ‘not-self’; or ‘God’ – as an interpretation of ‘universal being’). If so, we can ask what comes first. We might see this kind of orientation as intrinsically a moral issue: selfishness bad; unconditional generosity good. But that would suggest the source of moral evaluation is the moral faculty, inviting Molière’s jibe about opium putting you to sleep on account of its soporific virtue.²¹ But if the other way round, with the ‘heart’ as the real source of moral evaluation without circularity, can we understand this without assuming a ‘universal being’?

²⁰ ‘Technical rules’ presumably include CP.

²¹ Vertus dormitiva: see Molière (2003, 2012: Third Interlude).

With God in the equation, this seems to root morality in godliness. We can compare a self which loves God with one which loves itself, and see goodness only in the former. But a whiff of circularity remains. 'God' comes so packed with moral content it is hard to see how a 'heart' choosing to 'love God' is not, by its 'allegiance', intrinsically choosing to be good. What, though, if we remove God? Something's 'worth' is not necessarily a moral worth. It could be simply its value to oneself on any level – even just the fact that one wants it. James needs this 'heart' to be antecedent to moral evaluation, a faculty which is not 'science' or 'reason', can do moral evaluation ('compare the worths'), but can also 'love the universal being', 'harden itself against itself' and 'perceive God'.

James is three-quarters through his essay, having already explored the role of the 'will' (or 'willing nature') in belief-formation.²² So it is tempting to read 'heart' as not the 'will', otherwise James would have said 'will'. But despite this 'heart' and 'will' do seem very close – particularly as he continues:

The question of having moral beliefs at all or not having them is decided by our will. Are our moral preferences true or false, or are they only odd biological phenomena, making things good or bad for us, but in themselves indifferent? How can your pure intellect decide? If your heart does not want a world of moral reality, your head will assuredly never make you believe in one. (1896, 2000:212-3)

James seems to be saying that for an individual to have moral beliefs at all presupposes a decision to have moral beliefs.²³ This does not have to mean every individual with a moral belief has previously deliberately committed to having moral beliefs, but that having a moral belief involves committing to what having a moral belief entails. It would be possible not to commit, and instead remain outside the 'world' of morality.

This seems plausible, which is not to claim everyone, or every moral being, must

²² Covered in **Chapter 4**, especially p140ff.

²³ '[T]here are no actual moral truths or laws ... until we choose to play the [moral language-]game' (Gale, 1999, 2007:47)

agree. A moral realist, in particular a naïve moral realist, might counter that moral principles are objective, universally binding and inescapable; that personal commitment to a 'moral world' is possible but irrelevant to 'moral reality'. Someone with beliefs aligned to those objective, binding and inescapable moral principles would believe things which are true, regardless of whether she has made any commitment. But not everyone would subscribe to this either, and James's position (as interpreted here) seems no less defensible.

His coverage of moral belief (outside any moral implications of religious belief) concludes:

Moral scepticism can no more be refuted or proved by logic than intellectual scepticism can. When we stick to it that there is truth (be it of either kind), we do so with our whole nature, and resolve to stand or fall by the results. The sceptic with his whole nature adopts the doubting attitude; but which of us is the wiser, Omniscience only knows. (1896, 2000:213)

At first sight the parallel between moral and intellectual scepticism seems unconvincing. In the purely 'intellectual' domain the choice between embracing and not embracing scepticism may appear a 'forced option' (1896, 2000:199)²⁴ between two mutually exclusive and jointly exhaustive positions. But to assume an equivalent forced option holds in the moral domain implies only two 'hypotheses' apply:

- (2.7.1) There is a moral truth: therefore our attempts to behave morally are attempts to align ourselves with and/or acknowledge that moral truth.
- (2.7.2) There is no moral truth: therefore we have no reason to be good or behave morally.

But other alternatives exist, including:

²⁴ See p142ff.

(2.7.3) The commitment to behave morally and/or see human behaviour (including our own) in moral terms is a commitment we make, and we give the commitment importance by making it.

(2.7.3) neither proposes nor denies a moral 'truth' – so is distinct from both (2.7.1) and (2.7.2). It is also not vacuous, as an individual may not commit. But it could be argued that an equivalent to (2.7.3) can also apply in the 'intellectual' domain. Someone may decide to behave as if the world contained real objects and/or other minds independent of her. However this would seem to be an act of will, not another intellectual option.

We cannot progress more without pre-empting much of **Chapter 4**. For now we will suggest the resolution could depend on what exactly James means by both 'scepticism' and 'stick to it'.

This in turn means James is little help so far. He would exclude prescriptive beliefs from a principle like CP, but he might exclude other beliefs which we would want to include, because he would see little systematic difference between descriptive and (morally) prescriptive beliefs in terms of their susceptibility to evidence. My conviction that I am justified in believing you are taking drugs (1.3) because I have sufficient evidence presupposes rejecting scepticism. But if my rejecting scepticism is itself unjustified then all necessary links and continuities – between, say, the experiences which I believe constitute my evidence and what I take as my current experience of you – collapse. If moral scepticism and intellectual scepticism are equally incapable of being 'refuted or proved by logic' then I am equally justified, and equally unjustified, in believing (morally) prescriptive beliefs ('moral preferences') can be 'true or false' (and therefore in principle capable of being supported by evidence) and in rejecting intellectual scepticism – which I would have to do to consider any descriptive belief adequately supported by evidence.

2.6 Problems with Clifford's principle applied to prescriptive beliefs

We now return to the two rival morally prescriptive beliefs in Clifford's beggar story.²⁵

(2.5) We ought to give money indiscriminately to beggars.

(2.6) We ought to help people do the work they are most fitted for, and not keep and encourage them in idleness.

Clifford favours (2.6) over (2.5) on the grounds that (2.6) is now supported by sufficient evidence (the 'investigation') whereas (2.5) is relatively unsupported. The question is still what would count as evidence for a prescriptive belief.²⁶

For present purposes we will assume a statement, *s*, counts as evidence for a belief that *p* if *s* is true and *s* provides objective reasons for thinking *p* is true or is likely to be true.²⁷

Consider what might count as evidence for or against the descriptive belief that the moon is made of green cheese. *For*: From the earth the moon resembles Sage Derby. *Against*: Spectrographic analysis of lunar light gives results more like reflections off rock than off cheese; Apollo 11 landed somewhere hard. This evidence is publicly available, at least in principle. It is also expressed in publicly understandable, publicly agreed and objective terms. There is little debate about what 'cheese', 'spectrum' and 'manned spacecraft' refer to. It should therefore be straightforward to understand how statements describing spectrographic results and Apollo 11 discoveries could be objectively expressed and determined to be true or false.²⁸

²⁵ See p55ff.

²⁶ See p46ff.

²⁷ We will return to this working definition: see p215.

²⁸ See however **Appendix A3** p285.

(We should add a caveat. The concept of scientific evidence might be less clear-cut than this implies. Particularly during scientific revolutions there may be disagreement about what counts as evidence and how evidence bears on theoretical claims. Scientific revolutions themselves could present a challenge for an evidence principle. These issues are outside our current scope.²⁹)

Now consider what might count as equivalent evidence for prescriptive belief (2.6). There seem at least five levels of difficulty.

First level

We will make a utilitarian assumption that belief (2.6) is equivalent to the belief that:

(2.6.u) Helping people do the work they are most fitted for causes more net happiness than keeping and encouraging them in idleness.

It might be possible in principle to quantify total net happiness (say, balance of pleasure over pain) resulting from each strategy. This should cover all people (or all sentient beings) currently living or who will ever live, and who would qualify as beneficiaries or implementers of, or as being in any way touched by, each strategy.

The first level relates to whether quantification is practically achievable, not just in principle. If it is practically unachievable the evidence to support belief (2.6.u) would be unavailable – and therefore insufficient. Belief (2.6.u) would then fail CP. CP does not just say it is wrong to believe anything which could not in principle be supported by evidence, but that it is wrong to believe anything which is not in fact sufficiently supported by evidence, even if in principle it could be.

But consider a starker contrast:

²⁹ See **Appendix A5** p293.

(2.8) We ought to help people whenever we can and act out of sympathy for their suffering, rather than do whatever we can to hurt or destroy them.

Our utilitarian assumption makes (2.8) equivalent to:

(2.8.u) Helping people whenever we can and acting out of sympathy for their suffering causes more net happiness than doing whatever we can to hurt or destroy them.

Comparing resultant net happiness for the two strategies in (2.8.u) is almost trivial. It may still be practically impossible to calculate each total, but we only need relative sizes. We could be as confident in holding belief (2.8.u) on sufficient evidence as we would in believing every dropped object will fall to the ground because every object dropped so far has.³⁰

Our initial verdict on (2.6.u) was over-hasty. We only need enough evidence to support a comparison. If it is harder to find evidence for (2.6.u) than for (2.8.u) this could just indicate that belief (2.6.u) is less substantiated than (2.8.u), so we would have less right to believe it. It seems we could have sufficient evidence for (2.6.u) – although we may have a harder escape route for a more open-ended utilitarian belief:

(2.4.u) Being beneficent maximises net happiness.

Second level

The second level of difficulty concerns the move from descriptive belief (2.6.u) to (categorical) prescriptive belief (2.6), which our utilitarian assumption underwrote.

We have not yet seen how a (categorical) prescriptive belief could have evidence for or against it. If evidence is something which is true and provides objective reasons for thinking a belief is true or likely to be true, then we must be able to tell when the evidence and the belief are true.

³⁰ But see p95ff.

We will assume relevant evidence statements (about levels of happiness experienced by sentient beings when people are helped do the work they are most fitted for – assuming we can identify this – and about levels of happiness experienced by sentient beings when people are kept and encouraged in idleness) can be identified as true when they are true. We will also assume the relevant descriptive belief statement (helping people do work they are most fitted for causes more overall happiness than keeping and encouraging them in idleness) can be identified as true when it is true. So we can see how the truth of the evidence statements supports the truth of the descriptive belief statement.

It should be unproblematic to supply appropriate additional premises and/or definitions for descriptive belief statement (2.6.u) to entail an evidence statement like:

(2.6.u.e1) Levels of happiness observed in an appropriate context as experienced by sentient beings when people are helped do the work they are most fitted for are higher than levels of happiness observed in the same context as experienced by sentient beings when people are kept and encouraged in idleness.

Statement (2.6.u) entails one or more evidence statements. If these turn out true they would support (although not prove) the truth of (2.6.u). If the evidence statements were false they would falsify (2.6.u). So the converse belief statement

(2.6.u.2) Helping people do the work they are most fitted for causes no more overall happiness than keeping and encouraging them in idleness

would, with appropriate additional premises, entail an evidence statement like:

(2.6.u.e2) Levels of happiness observed in an appropriate context as experienced by sentient beings when people are helped do the work they are most fitted for are no higher than levels of happiness

observed in the same context as experienced by sentient beings when people are kept and encouraged in idleness.

Finding evidence statement (2.6.u.e2) false would entail that belief statement (2.6.u.2) is false. Which, in the current context, is close if not equivalent to entailing that original belief statement (2.6.u) is true.³¹

So we can understand how evidence (the truth of one or more evidence statements) can support descriptive beliefs (the truth of descriptive belief statements). But if we replace the descriptive belief with a prescriptive belief the relationship is less clear.

If the relationship were equivalent it would mean our (prescriptive) belief statement (2.6) would entail (descriptive) evidence statements of an equivalent form to (2.6.u.e1). But it does not, because an 'ought' statement like (2.6) can only entail another 'ought' statement like:

(2.6.b.1) If Fred is most fitted for work of type X, then we ought to help Fred do work of type X, and not keep and encourage him in idleness.

It is unclear how (2.6.b.1) can function as an evidence statement objectively verifiable as true or false.

The second level of difficulty is therefore that applying CP to prescriptive beliefs needs a solution to the 'is/ought problem'. We would first have to show the prescriptive belief is equivalent to one or more descriptive beliefs, or at least that there is enough conceptual connection between the two to allow evidence for the descriptive belief(s) to count as evidence for the prescriptive belief. If we could do this we could then ask if those descriptive beliefs are sufficiently supported by evidence. But that is a big if.

³¹ The context here is that both (2.6.u.2) and (2.6.u) are expressed as comparisons, and the equivalence applies because of the way they are worded. We are not making a more general claim that, for example. If putative evidence p for A turns out to be false, then it is evidence for not- A .

We are not claiming to have proven that CP cannot apply to (categorical) prescriptive beliefs. But in this subsection we have tried and failed to find a way to reduce prescriptive statements to descriptive statements. We are not claiming it is impossible, just that it is hard to see how it can be done, or that it can be done. Caution would therefore favour excluding prescriptive beliefs until we have a coherent account of how evidence supports them.

A related issue is that a utilitarian justification for preferring one prescriptive belief to another by comparing the weight of supporting evidence requires not only accepting in principle a utilitarian transition from 'is' to 'ought', but also agreeing the appropriate domain the utilitarian calculus applies to: all actual people; all actual and potential people; all actual and potential sentient beings; etc.

It could be countered that for a utilitarian or similar moral consequentialist prescriptive beliefs like (2.6) and (2.4) are formally equivalent to the kind of prudential belief which we established³² could be supported by evidence:

(2.1.1) You should fit the handle into the hole in the hatchet rather than tie the two together.

But identifying the right kind of evidence required treating (2.1.1) as a hypothetical belief with an assumed antecedent '[If you want to chop wood with less effort and less risk of injury then]...' (2.1.2). This hypothetical prescriptive belief was equivalent to a descriptive belief which could be supported by relevant evidence. But a different assumed antecedent³³ would need different evidence.

We could try the same approach with (2.6):

(2.6.1) **[If we want to maximise net happiness then]** we ought to help people do the work they are most fitted for, and not keep and encourage them in idleness.

³² p44ff.

³³ See (2.1.4) p45.

This is equivalent to a descriptive belief:

- (2.6.1d) Helping people do the work they are most fitted for, rather than keeping and encouraging them in idleness, will maximise net happiness.

In principle we could identify evidence supporting (2.6.1d), and therefore (2.6.1). But we could assume a different antecedent:

- (2.6.2) **[If we want to recreate 1940s social conditions then]** we ought to help people do the work they are most fitted for, and not keep and encourage them in idleness.

The equivalent descriptive belief might be:

- (2.6.2d) Helping people do the work they are most fitted for, rather than keeping and encouraging them in idleness, will recreate 1940s social conditions.

(2.6.1d) and (2.6.1) call for different evidence to (2.6.2d) and (2.6.2).

A utilitarian who believes (2.6) may also believe (2.6.1); and may believe (2.6) because she also believes (2.6.1). But believing (2.6.1) can only be a reason for believing (2.6) if she also believes the antecedent of (2.6.1): that we should maximise net happiness. We pursue this in the next subsection.

Third level

We suggested an answer to the 'is/ought problem' could be a kind of utilitarianism. One utilitarian might say 'good' or 'right' just *means* 'productive of the greatest net happiness' – which is vulnerable to Moore's 'open question' objection (1903, 1922:§13). If 'good' or 'right' just means 'productive of the greatest net happiness', then (assuming Moore is right) the question 'Is it true that producing the greatest net happiness is good?' would be either meaningless or trivial.

What Moore said (or is said to have said) on this subject has generated considerable debate. Feldman sees Moore as insisting that

you can't derive substantive conclusions in axiology from claims about the synonymy of 'good' with some other expression (especially if the other expression was either naturalistic or complex). (2005:40)

In Moore's words:

if I am right, then nobody can foist upon us such an axiom as that 'Pleasure is the only good' or that 'The good is the desired' on the pretence that this is 'the very meaning of the word.' (1903, 1922:7)

Adapting one of Feldman's 'Synonymy Argument' (2005:39) schemas:

- (2.9.1) 'x is good' means 'x is productive of the greatest net happiness'.
- (2.9.2) If 'x is good' means 'x is productive of the greatest net happiness', then something is good if and only if it is productive of the greatest net happiness.
- (2.9.3) If something is good if and only if it is productive of the greatest net happiness, then utilitarianism is true.
- (2.9.4) Therefore, utilitarianism is true.

Feldman argues Moore was not claiming (in section 13 of *Principia Ethica*) the consequent of (2.9.2) is false; but that premise (2.9.1) is false and therefore conclusion (2.9.4) is not proven. Even if (2.9.1) is false, the consequent of premise (2.9.2) could still be true. 'Something is water if and only if it is H₂O' can be true even though 'water' and 'H₂O' are not synonymous. This seems sound. 'Is it true that producing the greatest net happiness is good?' and 'Are water and H₂O the

same thing?’ are not questions about word meanings like ‘Is a bachelor an unmarried man?’

But another utilitarian approach might express itself as a moral principle: not that ‘good’ or ‘right’ means ‘productive of the greatest net happiness’, but that we (morally) ought to act to maximise net happiness. The consequent of premise (2.9.2) would then hold, but on the assumption that saying an action or choice is ‘good’ is a morally prescriptive claim.

This approach avoids the ‘naturalistic fallacy’ and legitimises the move from ‘is’ (2.6.u) to ‘ought’ (2.6) by supplying missing premise (2.6.m):

(2.6.m) We ought to do what causes more net happiness.

(2.6.u) Helping people do the work they are most fitted for causes more net happiness than keeping and encouraging them in idleness.

Therefore:

(2.6) We ought to help people do the work they are most fitted for, and not keep and encourage them in idleness.

But (2.6.m) is prescriptive. If CP covers prescriptive beliefs, it applies to (2.6.m). It is therefore only right to believe (2.6.m) on sufficient evidence. But we do not yet understand how evidence can support a prescriptive belief, without one or more additional premises to get us to an equivalent descriptive belief which evidence can support. This seems an infinite regress.

But have we exhausted utilitarian approaches? We have considered ‘synonymy’ (2.9.1) and ‘moral principle’ (2.6.m). Is there a variant of (2.6.m) which is not itself prescriptive? This would also claim something is good (and therefore what we morally ought to do) if and only if it maximises net happiness, but as a descriptive, not a prescriptive claim. It would be a claim about the universe which, if true,

would sanction the move from an appropriate 'is' statement to a corresponding 'ought' statement:

(2.6.d) It is a fact about the universe that something is our duty if and only if it causes more net happiness than any other option.

But even if (2.6.d) were tenable as a descriptive belief, it would still be subject to CP. It would only be right to believe (2.6.d) on sufficient evidence. Again it is hard to see what could count as evidence for (2.6.d): but not, ostensibly, because of the infinite regress we had with (2.6.m). (2.6.d) is by definition a descriptive belief, assuming it can be understood as such.

Someone could perhaps claim that a reason for *not* believing (2.6.d) is that there are actions which cause more net happiness but which are clearly not good. A reason for believing (2.6.d) is false would be evidence against (2.6.d). Conversely an absence of such actions would be evidence supporting (2.6.d). But if that was the case then the very same actions would constitute evidence against prescriptive belief (2.6.m): the evidence would be that (2.6.m) leads to endorsing as good actions which are clearly not good. The problem is unpacking 'clearly not good'. Clear to whom? Surely only to someone holding one or more prescriptive beliefs incompatible with (2.6.m). Those beliefs – '(2.6.not-m)' – would be why she saw the actions as evidence against (2.6.d) or (2.6.m) or both. They would only count as evidence if (2.6.not-m) is true. What though would count as evidence for (2.6.not-m)? Not the same actions claimed as evidence against (2.6.d) or (2.6.m), as these presuppose the truth of (2.6.not-m). But if the evidence consists of other actions then these would only count as evidence if they are identifiable as good or bad (or neutral) from the viewpoint of one or more prescriptive beliefs – which may be, or include, (2.6.not-m). This seems another endless search. We will take the apparent unattainability of evidential support for (2.6.d) to indicate it is not a serious contender.

The next level of difficulty relates to the implications of a principle like CP applying to itself when held as a belief.

Fourth level

CP says it is wrong to believe anything on insufficient evidence. If CP covers prescriptive beliefs, then it would be wrong to believe CP on insufficient evidence. But we have not yet established how, or indeed whether, evidence can support a prescriptive belief. Until it is established that there can be sufficient evidence for CP, the belief that CP ('bCP') would invalidate itself. This problem does not arise if CP only applies to descriptive beliefs, because then it would not apply to bCP.

Fifth level

Consider the predicament faced by an agent deciding whether to adopt CP. We will assume we have not yet ruled out the possibility of evidence for prescriptive beliefs in general or, in particular, for bCP.

We start with the agent in a state of indifference (S_i), neither accepting nor rejecting bCP. A body of potential evidence E appears. Because of her indifference she has no reason to apply CP to bCP, therefore no reason to consider whether E supports CP. She could stay in state S_i indefinitely. Or she could move to assumed acceptance of bCP (S_a) or assumed rejection of bCP (S_r). In state S_a she would have reason to consider whether E supports bCP, because her assumed acceptance of bCP would incline her to do this with any of her potential beliefs. She might find E does support bCP, which would then lead her to confirm her acceptance of bCP.

But she could have moved to S_r , where she would have no reason to link E with bCP, so bCP would stay outside her set of beliefs.

The issue is: why would she 'bootstrap' herself into believing CP, and what would justify her in doing this? Even assuming evidence could exist which supports bCP it seems she would only entertain bCP by presupposing its truth, which CP itself would stop her doing.³⁴

³⁴ Imagine a universe ruled by a God who has ordained a set of prescriptive rules which the inhabitants must obey. One of the rules is CP. The inhabitants know the God has ordained CP, but

2.7 Conclusions

We have not proved CP cannot apply to prescriptive beliefs, which was our stronger potential conclusion.³⁵ But significant issues arise when applying CP to prescriptive beliefs which do not arise with descriptive beliefs. This suggests a principle (EP1) excluding prescriptive beliefs would be more defensible than CP, and that excluding prescriptive beliefs is neither unsound nor arbitrary. This weaker conclusion is strong enough for our purposes.

We are after all assuming EP1 is a moral principle which people could feasibly apply to their behaviour. If 'ought' implies 'can' it would seem unreasonable to expect people to support their prescriptive beliefs with evidence if we cannot explain how this is possible.

also know its wording compels them to accept bCP only on sufficient evidence of its truth, not because the God has ordained it. But CP would forbid them applying bCP until they knew it was true by having sufficient evidence of its truth, and therefore of the truth of CP.

³⁵ p43.

Chapter 3 | Against an evidence principle

Our principle is currently:

EP1 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything on insufficient evidence.

This chapter will look for reasons for thinking EP1 cannot be right. It will consider apparent counter-examples, to identify instances of descriptive belief which might lead us to modify or rethink EP1 – perhaps by incorporating additional inclusions, exclusions or caveats – or help clarify the intention behind EP1. Its primary objective is not so much to build a positive case for EP1, as to see where more nuance may be needed and formulate sufficiently robust wording.

3.1 Involuntary beliefs

EP1 declares over-belief within its scope morally wrong. If this includes involuntary beliefs (assuming they are possible) we might be claiming it is morally wrong to acquire or hold an over-belief even if acquiring or holding it was out of our control.

This seems counter-intuitive. Can it be morally wrong to say something untrue in your sleep, or hurt someone accidentally by tripping on a kerbstone? Nothing in Clifford's essay suggests he would say yes, although he discusses voluntary acts and omissions which encourage or enable involuntary over-belief in oneself and others. The ship-owner 'had knowingly and willingly worked himself' into a state where he 'felt so sure about it that he could not think otherwise' (1877, 1879:178).

Indeed concern about varying degrees of voluntary control (including no voluntary control) over belief could be why Clifford outlaws ostensibly harmless over-beliefs:¹ because they can 'weaken our powers of self-control, of doubting, of judicially and fairly weighing evidence' and replace 'the habit of testing things and

¹ See p121ff.

inquiring into them' with one of 'believing for unworthy reasons' (1877, 1879:185-6). Our self-control could dissipate, so we need vigilance over aspects of belief which are under our control, while they are.²

We said earlier³ that, other things being equal, the believer would be morally wrong to the extent that he acquired or held the descriptive over-belief voluntarily. We should make this explicit:

EP2 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything on insufficient evidence **[except when the unjustified believing is outside the believer's voluntary control]**.

The addition is to remove doubt. When the moral domain is explicit excluding involuntary acts should be taken as read. However the clause applies whether we want to highlight the belief itself or its being an *over*-belief as involuntary: 'voluntary' and 'involuntary' can apply at more than one point. EP2 would for example correctly exonerate over-belief resulting from cognitive inadequacy, to counter one of Haack's (1997, 2001:25) worries about CP.⁴

We are assuming involuntary beliefs are possible. Williams however doubts whether *voluntary* beliefs are possible: 'if I could acquire a belief at will, I could acquire it whether it was true or not' (1973:148). This would be problematic: if voluntary beliefs are impossible and involuntary beliefs are excluded, why have a principle which no beliefs fell under? Williams's 'Conceptual Impossibility Thesis' (Feldman, 2001:79) has been critiqued by a number of writers (eg Winters (1979)), and seems to have little current support. However what Feldman calls Alston's (1988) 'Contingent Inability Thesis' – 'that as a contingent matter of fact, people are not able to acquire beliefs voluntarily' (Feldman, 2001:80) could be an equivalent challenge.

² See **Voluntary and involuntary beliefs** p17ff.

³ p21 and p34.

⁴ p32.

To neutralise these worries we must distinguish direct from indirect voluntary control. Following Clifford we are targeting voluntary aspects of belief acquisition, maintenance and rejection. ‘Doubting’, ‘judicially and fairly weighing evidence’, ‘testing’ and ‘inquiring into’ our beliefs: these we decide to do or not do or remember or forget to do. We can generally exert direct control over these, even if they only indirectly control beliefs arising from or influenced by them.

Should we then change ‘believe anything on insufficient evidence’ in EP2 to refer strictly to those voluntary behaviours which indirectly lead to over-belief? That seems gratuitous. A moral principle against causing unnecessary pain would still be breached by someone doing (or failing to do) something which (or the omission of which) only indirectly led to unnecessary pain, if he knew or should have known this was likely, whatever his intention. The principle aims to avoid unnecessary pain, making sense as a behavioural guideline even in, say, a medical context where many potential acts and omissions likely to influence whether someone suffers unnecessary pain operate only indirectly.

We can explore this more systematically using Alston’s four ‘modes of voluntary control’ (1988:278):⁵

I. Direct control.

A. Basic control.

B. Non-basic immediate control.

II. Long-range control.

III. Indirect influence.

These represent different ‘strengths’ of doxastic voluntarism: one voluntarist might claim we have direct basic control (IA) over our beliefs; another that we have direct non-basic but immediate control (IB) over our beliefs; and so on.

IA Direct basic control. The model here would be

⁵ Later analytic approaches are available. McHugh (2014) for example models doxastic freedom on freedom of intention rather than freedom of action, and argues that ‘if we had voluntary control of ... doxastic states, this would actually undermine our freedom’ (2014:1). Alston’s analysis however is adequate for our purposes here.

the maximally direct control we have over the motions of our limbs and other parts of our body, the voluntary movements of which constitute “basic actions”, ... things we “just do”, not “by” doing something else voluntarily. (1988:260)

Few if any doxastic voluntarists think we can ‘believe at will’ like this.

IB Non-basic immediate control. An example would be deliberately turning a light on, where success requires ‘one or more bodily movements’ and ‘more than a volition’ on my part (1988:269). Despite this I might be blamed for not turning on the light if I should have done.

This also seems unlikely to apply to belief in any relevant way. I could voluntarily choose to gather evidence and as a result acquire a belief which that evidence would support. But at most that would indicate ‘immediate voluntary control over whether I take up *some* propositional attitude [like belief] toward *some* proposition’ (1988:271) relating to the evidence gathered. We can also ‘control our evidence and thereby control our beliefs’ (Feldman, 2001:83), for example by switching the lights on in order to believe the lights are on (2001:82). But this is also irrelevant to EP2.

II Long-range control. This is ‘the capacity to bring about a state of affairs, C, by doing something (usually a number of different things) repeatedly over a considerable period of time’ (Alston, 1988:275). C could be losing weight or learning a language – or at least some of our beliefs:

[P]eople do set out on long range projects to get themselves to believe a certain proposition, and sometimes they succeed ... [for example by] selective exposure to evidence, selective attention to supporting considerations, seeking the company of believers and avoiding non-believers, self-suggestion, and (possibly) more bizarre methods like hypnotism. (1988:275).

Interestingly Alston does not see this as 'a kind of voluntary control that grounds deontological treatment' – treatment like EP2 – because

people could properly be held responsible for their attitudes toward propositions in a certain range only if those who set out to intentionally produce a certain attitude toward such a proposition, and made sufficient efforts, were frequently successful. ... [Otherwise, even] if I had done everything I could to produce [a certain result], I would have had little chance of success; so how could I rightly be blamed for its absence? (1988:275-6)

Alston doubts that we have reliable long-range control over any of our beliefs, even religious or philosophical beliefs or beliefs about personal relationships. People may 'sometimes' succeed in getting themselves to believe or disbelieve something. But the success rate would not be 'substantial'.

We must be clear what this category is. It is where the agent 'manages to believe something contrary to what seems to him to be the case or something concerning which he has no definite impression of truth or falsity' (1988:277). It is not where an agent forms or retains beliefs 'without adequate grounds, reasons, or justification': where the belief itself would typically not be voluntary, because the relevant proposition 'seems clearly true, however ill supported' (for example prejudice where, say, 'socialization' has led to it seeming 'clearly true' that one race is superior to another). Category II is rather where the agent intends to believe that p , and succeeds in believing that p , by one or more voluntary actions or behaviours which execute that intention.

It is not crucial to our case to decide whether this category II is conceptually impossible, contingently impossible, or just sparsely populated. What matters is that Alston's first three categories, which treat 'propositional attitude formation on the model of intentional action' (1988:277), collectively fail to identify most of the beliefs and doxastic behaviours which EP2 targets.

III Indirect influence. We can also be held responsible for a state of affairs even if we did not produce it intentionally, provided something we did but should have not

done (or did not do but should have done) was 'a necessary condition (in the circumstances) of the realisation of that state of affairs' (1988:278): provided, that is, the state of affairs would not have resulted had we not done something we should not have done (or done something we should have done). For example I may not be able to will myself to understand (or not understand) something, but if I deliberately choose not to read the relevant textbooks I am still responsible for not understanding the causes of WW1, if I (probably) would have done had I done my history homework.⁶

Alston considers even this category provides no grounds for 'a deontological conception of epistemic justification in terms of freedom from blame in taking up a certain propositional attitude', and concludes we are 'ill advised to think of epistemic justification in terms of freedom from blame for believing' (1988:294). However EP2 is not intended as a comprehensive principle of epistemic justification derived from moral considerations, like for example Haack's 'correlation thesis',⁷ but as a self-standing moral principle applying to doxastic behaviour. For our purposes Alston's category III indirect influences are the principal acts and omissions in EP2's intended scope.

We will now test this interpretation of where volition applies in acquiring and holding beliefs. First a belief resulting from brainwashing, assuming this is possible psychologically. By definition it would be acquired on insufficient evidence. Would its acquisition be immoral? If the brainwashing was involuntary, the belief would fall outside EP2, and possibly also CP.

But what of voluntary hypnosis, to see if it was possible to acquire an over-belief involuntarily? As EP2 is currently worded it would be disallowed. Then if, having acquired through hypnosis a belief in (say) the existence of a supernatural being, the agent becomes aware of having the belief, EP2 would oblige him to test the belief against evidence and reject it if it fails. If there was something about the belief (or the way it was acquired) which made it impossible for him to become

⁶ We will have more to say about indirect influences on belief later in the context of Pascal's Wager: p145ff.

⁷ This is that positive or negative epistemic appraisal is 'invariably' and 'contingently' associated with positive or negative moral appraisal (1997, 2001:21): see p31ff.

aware of having the belief or to apply EP2 to it, then under EP2 he could correctly be blamed for voluntarily allowing this to happen. If this seems severe consider someone deliberately undergoing hypnosis to become irreversibly indoctrinated in Nazi ideology.

Now for a subject undergoing voluntary hypnosis in pursuit of psychological enquiry, where (say) the desirable outcome of enhanced knowledge was considered more likely than the undesirable outcome of a fixed over-belief. As EP2 is a moral principle it can be assumed to include the implication that any particular instance of over-belief would only be morally wrong in the absence of any conflicting and overriding moral imperative⁸ – which in this case would be to further scientific knowledge. We will make this explicit:

EP3 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything on insufficient evidence **{in the absence of any conflicting and overriding moral imperative}** [except when the unjustified believing is outside the believer's voluntary control].

This latest contextual presupposition⁹ should again be technically redundant as a potential context of conflicting moral imperatives should be assumed in the moral domain. EP3 is not claiming itself superior to any other moral principle, about which exactly the same could be said.

The 'conflicting and overriding moral imperative' refers to the specific instance of (for example) actual or potential behaviour, not to the generic moral principle the instance may fall under. This is another commonplace of moral choice: an instance of behaviour could qualify as, say, both 'telling the truth' and 'avoidably hurting an innocent person', which is how moral conflict arises. The dilemma is in deciding whether in *this* particular instance the imperative to tell *this* truth overrides the imperative not to hurt *this* innocent person, or vice versa. EP3 has no privileged position in a hierarchy of principles.

⁸ See p22.

⁹ See p22.

The broader question of voluntary versus involuntary aspects of belief therefore seems unproblematic for EP3. EP3 implicitly includes voluntary behaviour and explicitly excludes involuntary behaviour, relating to acquiring and/or holding beliefs. The inclusion and exclusion do not rob EP3 of its intended scope, and are not arbitrary, because the specified wrongness is moral wrongness.

3.2 Foundational beliefs

Some of the foregoing discussion will also apply to so-called ‘foundational’ or ‘basic’ beliefs.¹⁰ If a belief is supported by evidence, that evidence would typically consist of beliefs, which in turn rest on other evidence, also typically in the form of beliefs. The chain must presumably end somewhere – perhaps in ‘foundational’ beliefs unsupported by other beliefs? Plantinga speaks of an

apparent cleavage between those beliefs you accept on the evidential basis of other beliefs, and those you accept in the basic way—accept, but not on the evidential basis of other beliefs (1993:68)

Would these basic beliefs be evidentially unsupported, and therefore over-beliefs? Such beliefs are considered problematic for some kinds of evidentialism, so we must assess their implications for EP3.

Exploring the status and role of foundational beliefs has led to a range of views known as ‘foundationalism’. Plantinga is critical of foundationalism, but he offers a useful survey of candidate types of foundational beliefs, along with a terminology. For example ‘warrant’ is what ‘makes the difference between knowledge and mere true belief’ (1993:3):

To say that a belief is warranted or justified for a person is to evaluate it or him (or both) positively; his holding that belief in his circumstances is right, or proper, or acceptable, or approvable, or up to standard.

¹⁰ We will use ‘foundational’ and ‘basic’ interchangeably.

Warrant is not however the same as justification (1993:46). For Plantinga a belief can be warranted without being evidentially supported. For example, both basic beliefs and 'superstructure' beliefs (beliefs derived from and/or supported by basic beliefs) can be warranted; and a warranted basic belief will count as 'properly basic':

We could hold that a belief is properly basic for me only if it has so much warrant that it is certain for me (has the maximal degree of warrant for me); or ... if it has enough warrant so that I know it (or have some special type of knowledge of it). Alternatively, we could say that it is properly basic for me if I am not irrational in accepting it, am justified in accepting it, am within my epistemic rights in accepting it... (1993:70)

Different 'foundationalists' define 'properly basic' beliefs in different ways. Plantinga thinks Locke for example holds that 'a belief—any belief—is properly basic for me only if it is either self-evident or appropriately about my own immediate experience' (1993:71). But for a superstructure belief to be warranted it would have to be derived (immediately or ultimately, perhaps through chains of linking beliefs) from one or more 'properly basic' beliefs. A basic belief which was not warranted (so not 'properly basic') could not support any warranted superstructure beliefs.

Plantinga (1981) and (1983) provide examples of beliefs which could be considered properly basic:

- (3.1.1) Simple, apparently self-evident, mathematical truths, eg ' $2 + 1 = 3$ ' (1981:41).
- (3.1.2) Simple, apparently self-evident, logical truths, eg 'No man is both married and unmarried' (1983:55).
- (3.1.3) Perceptual beliefs, eg 'I see a tree' (1981:47).

(3.1.4) Beliefs ascribing mental states to others, eg 'that person is pleased'.

(3.1.5) Memory beliefs, eg 'I had breakfast over an hour ago'.

All basic beliefs, therefore all properly basic beliefs, are relativised to the believer. So in one of Plantinga's proposed wordings for a 'fundamental principle of classical foundationalism':

(3.2) *A proposition p is properly basic for a person S if and only if p is either self-evident to S or incorrigible for S or evident to the senses for S.*
(1983:59)

This would apply not only to beliefs like (3.1.3–5) but also to 'self-evident' beliefs like (3.1.1–2). Plantinga sees properly basic beliefs as relativised both to the believer and to a set of conditions, which may or may not obtain:

Very simple arithmetical truths will be self-evident to nearly all of us, but a truth like $17 + 18 = 35$ may be self-evident only to some (1983:56).

[A] belief is properly basic only in certain conditions; [which] are ... the ground of its justification and, by extension, the ground of the belief itself.
(1983:80)

Beliefs like (3.1.3–5) though are all 'relatively specific and concrete propositions' (1981:47) rather than 'more general and abstract' propositions which they 'immediately and self-evidently entail', for example, respectively:

(3.3.3) There are trees.

(3.3.4) There are other persons – and indeed other minds.

(3.3.5) The world has existed for more than 5 minutes.

But we can also 'speak a bit loosely' and call (3.3.3–5) properly basic.

For our purposes the precise extent and definition of ‘properly basic’ beliefs matter less than identifying a category of beliefs which we apparently regard as justified (or warranted) but which cannot coherently be considered supported by evidence (in the form of other justified beliefs) without infinite regress or circularity, because our other beliefs depend on them for their justification or warrant.¹¹ Acknowledging such beliefs (which Plantinga accepts) does not mean committing to any flavour of ‘foundationalism’ (which Plantinga rejects).

Plantinga sees foundationalism as a ‘*normative* thesis ... about how a system of beliefs *ought* to be structured, about the properties of a correct, or acceptable, or rightly structured system of beliefs’ (1993:72). This includes mandating that a ‘proper [belief system] will have a foundation: a set of beliefs not accepted on the basis of other beliefs’ (1993:73). Plantinga accepts there are basic beliefs:

a proposition is in the foundations of my [belief system] if and only if it is basic for me, and it is basic for me if and only if I don’t accept it on the evidential basis of other propositions. This much of foundationalism should be uncontroversial and accepted by all. (2000:75)

However he denies the stipulations of foundationalism regarding the ‘conditions of proper basicity’, and therefore how a belief system ought to be structured. He rejects what he calls the ‘classical package: evidentialism, deontology, and classical foundationalism’ (2000:79). ‘Deontology’ refers to the prescriptivity of these stipulations: ‘duty, obligation, permission, being within your rights’ (2000:77). Evidentialism is

the view that belief ... is rationally justifiable or acceptable only if there is good evidence for it, where good evidence would be arguments from other propositions one knows (2000:64).

Here ‘rationally justifiable or acceptable’ expresses the deontological element, and

¹¹ But see Adler’s alternative approach p88ff.

the 'other propositions one knows' are either other evidentially supported beliefs or, ultimately, properly basic foundational beliefs – 'properly basic' being another normative concept.

Plantinga sees 'many forms of classical foundationalism' as 'self-referentially incoherent' because of deontological stipulations such that

a proposition A is acceptable for me if and only if it is either properly basic or believed on the evidential basis of propositions that are (1) properly basic, and (2) support A. But this proposition itself is not properly basic by this criterion: it is neither self-evident nor appropriately about someone's immediate experience, and ... it is certainly hard to see that it is appropriately supported by propositions that do meet that condition.
(1993:85)

If foundationalism collapses, so does evidentialism – if foundationalism posits properly basic beliefs as the foundation a believer's entire structure of evidentially supported beliefs depends on.

What Plantinga sees as 'uncontroversial and accepted by all' may not however ring true for Adler, who sees evidentialism's alleged vulnerability vis-à-vis foundationalism very differently: that evidentialism would be false not if foundationalism is *false* but if foundationalism is *true*:

were evidentialism to require foundationalism, evidentialism would be self-refuting, since the essence of foundationalism is that some beliefs are exempt from the demand for (noncircular) reason-backing. So the alleged dependency is really a way to deny evidentialism. (Adler, 2002:174)

For now though we will assume foundational beliefs are conceptually sound and consider if they are an issue for EP3.

One obvious defence would be that most of our foundational beliefs, most of the time, would be involuntary. We have said EP3 excludes involuntary behaviour

relating to acquiring and/or holding beliefs. This would typically apply both to 'specific and concrete' beliefs like (3.1.1–5) and to 'more general and abstract' beliefs like (3.3.3–5), and therefore other everyday metaphysical assumptions like the belief that nature is uniform and will continue to be so. We could also include what James calls 'common sense' (1907, 2000:74): those 'intellectual forms or categories of thought' (1907, 2000:77) (like 'thing', 'the same or different', 'kinds', 'minds', and 'causal influences'): at least insofar as we could see them as beliefs, or as giving rise to beliefs.¹²

To consider their implications for EP3, we should distinguish between (i) structures (categories, concepts) which, whatever their source, are part of the architecture of the mind and are preconditions for our kind of thinking; and (ii) metaphysical propositions believed and/or asserted as statements about the fabric of reality which (may) derive their content from those aspects of mental architecture (i). Take for example a potential foundational belief that causation is something real explaining instances of constant conjunction. This would be incompatible with an alternative belief that causation is an illusion resulting from observations of constant conjunction. Neither belief is supported by evidence, if the only available evidence is the constant conjunctions themselves. It might seem otiose to condemn, under EP3, a person's over-belief in (say) any reality to 'causation' over and above scientifically precise observations of constant conjunction, even where this belief has ceased to be involuntary. Yet if it is 'natural' (say, in the sense of being part of evolved human nature) to believe causation is something real explaining constant conjunction (rather than just a way of labelling constant conjunction) then it might be equally 'natural' to infer from patterns in nature to the existence of agency in rivers, volcanoes and thunderstorms, when closer observation and/or more careful evaluation of evidence might cast doubt on such a 'natural' inference.

For many people, for much of the time, many foundational beliefs would count as

¹² Peirce has a similar conception:

[N]ot man merely, but all animals derive by inheritance ... two classes of ideas which adapt them to their environment[:] ... notions ... of force, matter, space, and time; and ... notion[s] of what sort of objects their fellow-beings are, and of how they will act on given occasions.
(1955:214-5)

involuntary, and so fall outside EP3. But the more stubbornly a believer persists in a 'natural' belief (that, say, agency exists in natural phenomena) by refusing to countenance alternative explanations and/or analytical evaluation of the evidence, the less it would count as involuntary. Analytical evaluation of the evidence may sometimes only justify rejecting the belief, not replacing it with a different one. It could justify the position that belief either way is inappropriate. An over-belief is not justified just because its only competition is other over-beliefs.

Insisting on holding an over-belief because there is no counterevidence would qualify as 'believing for unworthy reasons' (Clifford, 1877, 1879:185) just as much as deliberately avoiding counterevidence. Nor is it obvious that a belief in the 'reality' of causation or of deliberate agency within natural phenomena would count, or would need to count, as a 'foundational' belief upon which a believer's other beliefs depend, unlike (say) beliefs in the existence of other minds or in the continued uniformity of nature.¹³

The last two examples bring us to a second reason why (genuine) foundational beliefs need not be problematic for EP3. We have said many foundational beliefs would count as involuntary, but not all would. After all, a belief involuntarily acquired and/or held could transition into one we consciously and deliberately insist on, despite sceptical challenges. Our second reason is that EP3 as a moral principle presupposes a social context to belief.

An immediate objection might be that some potential moral obligations, permissions and prohibitions just concern the individual: for example Kant's case of the man whose misfortunes have

mounted to the point of despair, but he is still so far in possession of his reason as to ask himself whether taking his own life may not be contrary to his duty to himself (2005:98).

But even if this is correct in general it is not relevant to EP3, as the principal

¹³ We are including belief in the uniformity of nature here under the general category of foundational beliefs, but will also discuss it more specifically in the next section: p90ff.

rationale behind EP3 as a moral principle is the social impact of either obeying or disobeying it.

This presupposed social context means EP3 cannot disallow potential foundational over-beliefs in, say, the continued uniformity of nature or the existence of other minds if this would undermine that social context. Social life would be unthinkable without shared social projects, which must presuppose nature's continued uniformity. EP3 would permit such 'socially necessary' foundational beliefs not only because of its implicit social context as a moral principle but also because of its explicit caveat 'in the absence of any conflicting and overriding moral imperative'.

Our third defence relates to Plantinga's complaint that foundationalism is 'self-referentially incoherent' (1993:85). Even if this holds for foundationalism there seems no reason to accept an equivalent argument against EP3, which is proposed not as a foundational belief itself but as a moral principle covering only descriptive beliefs. Holding EP3 would mean holding a prescriptive belief which does not apply to itself: therefore not 'self-referentially incoherent'.

These three defences presuppose what Plantinga takes as 'uncontroversial and accepted by all' (2000:75): that someone can hold basic beliefs not on the evidential basis of other beliefs but which are foundational to the rest of his belief system. Plantinga accepts this but rejects foundationalist prescriptions as to which basic beliefs count as 'properly basic' and can therefore be legitimate foundations. Adler takes the different line of 'denying that there are basic beliefs' (2002:163).

There is only space to sketch a possible fourth defence based on Adler's approach. He sees beliefs like (3.1.1–5) and (3.3.3–5)¹⁴ not as 'basic' beliefs which would be either 'self-justified' or 'justified but not by further beliefs' but as members of 'our vast collection of *well-founded background beliefs*'. Beliefs like these can still do the 'job ascribed to basic beliefs – to provide a foundation' – but they

¹⁴ And his own examples like 'ants do not study the calculus' and 'objects do not disappear when not perceived'.

cannot do the philosophical work of basic beliefs precisely because they are well founded, and well founded because massively confirmed. The confirmation is tacit, a by-product of many everyday day activities.

We should bear in mind Adler's insistence on the 'modest ambitions of the ethics of belief' (2002:17), for example 'by ruling out radical skeptical alternatives and their surrogates' (2002:7). So when he claims that

the existence of other minds is tacitly confirmed, the key implication is that you would not in ordinary argument and discussion criticize someone who implied that he believed that there were other minds. (2002:185)

An opponent could reject these 'modest ambitions' and challenge whether our background beliefs really do 'receive massive tacit confirmation' (2002:166). But tacit confirmation of background beliefs does not mean that, in Quine's words, 'no statement is immune from revision' (1980:43). Adler thinks Quine does not 'distinguish ... between falsification and revision' (2002:165). Any statement, therefore any belief, can be falsified. But

[r]evision, which presupposes falsification, enters the further claim that if we assign the value "false" (falsified) to a statement previously accepted (as true) then we can follow through with all related alterations of truth-value to yield a new, coherent, and simple system.

Such revision is 'not guaranteed to succeed':

If we falsify certain background beliefs like that the universe exists or that there is no all-powerful Evil Genius bent on deceiving us, which are empirical and contingent, we also might not know how to form a new, coherent, and simple system. For to assign "false" to beliefs that receive massive tacit confirmation is to undermine our reliance on any evidence or experience for inference. (2002:165-6)

It seems therefore that a defence of EP3 based on Adler's concept of 'well-founded background beliefs' as an alternative to foundational basic beliefs might also benefit from something like our second defence, depending on the presuppositions behind the envisaged attack. We could argue that social life would be untenable without 'reliance on any evidence or experience for inference'. But in this case the social context presupposed by EP3 as a moral principle would legitimise beliefs in the existence of (say) the external world or other minds not as foundational over-beliefs but as tacitly confirmed background beliefs. We should acknowledge though that this approach may not succeed with the belief in the continued (and therefore future) uniformity of nature, which by its nature could never be considered tacitly confirmed, as confirmation can only be retrospective, and therefore inconclusive.

In this section we have avoided committing to either 'foundationalist' or 'anti-foundationalist' positions, and instead considered to what extent 'basic' or 'foundational' (or in Adler's case 'background') beliefs might undermine EP3. Following our brief survey we conclude that, with the possible exception of belief in the continued uniformity of nature, we have no reason to think EP3 as a moral principle (presupposing a social context) should disallow these beliefs, whether seen as foundational or as tacitly confirmed background beliefs.

We will discuss belief in the continued uniformity of nature further in the next section.

3.3 Past and future; testimony; uniformity of nature

Belief in the uniformity of nature could correctly survive EP3 when it is involuntary, but it may not always be involuntary. It could also survive EP3 by qualifying as a 'foundational' belief. But it features in other categories of potentially problematic beliefs: those about the past and, more obviously, about the future. Much of Clifford's defence of CP consists of guidelines for avoiding excessive scepticism: on beliefs about the past (1877, 1879:208-9); about the future (1877, 1879:206-7); or on another's authority or testimony (1877, 1879:188-199). Many such beliefs directly or indirectly presuppose the uniformity of nature, which might seem

problematic for CP and therefore perhaps for EP3.

This section discusses dependencies within this cluster of notions, which will expand into the nature of belief in general, moral belief in particular, and moral life. We also foreshadow coverage in **Chapter 5** of the shared social asset of knowledge and true belief.¹⁵

Beliefs about the past

Beliefs about the past do not divide neatly into personal memories and beliefs supported by testimony, as these overlap. It would be hard to imagine a testimony-based belief involving no personal memory: otherwise how was the belief acquired? The reverse overlap may be looser: remembering going upstairs five minutes ago may involve no significant ‘testimony’; but if my memory of hearing Handel’s Messiah last year goes beyond mere phenomenology it would have to incorporate *some* trusted cultural knowledge about Handel’s life and works.

We have covered personal memories explicitly under foundational beliefs and implicitly under involuntary beliefs. So they should not be problematic for EP3. Beliefs based on testimony are another matter though. Few could be classed as either foundational or involuntary – the latter in the sense of beliefs we could not avoid holding or acquiring. Our earlier history homework story¹⁶ serves as a test case. The student assumes his books tell the truth about the causes of WW1, because he assumes their authors had good reason to believe their own original sources. But a sceptic might question whether the student had ‘sufficient evidence’ for accepting this chain of testimony.

Testimony

Clifford asks if we are justified in believing Thucydides’ account of the siege of Syracuse (1877, 1879:208). And perhaps

¹⁵ For example p210ff.

¹⁶ p79.

(1.3) I believe you are taking recreational drugs¹⁷

because I was told this by person A, who I had heard was not exactly trustworthy. So I was unjustified in believing A's testimony, and therefore unjustified in belief (1.3). But what if I had been misled about A, who had been telling the truth all along?

Despite this layered testimony (1.3) still seems a case of over-belief (in the context sketched here), even if it turned out true. From the information I had at the time I should not have believed (1.3). We will incorporate this nuance:

EP4 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything [**knowingly or irresponsibly**] on insufficient evidence {in the absence of any conflicting and overriding moral imperative} [except when the unjustified believing is outside the believer's voluntary control].

The '[knowingly or irresponsibly]' includes cases where we are, or should have been, aware, or had good reason to believe, that our evidence was non-existent, insufficient or compromised, either in itself or in how we are taking it. This is not unique to testimony but it is clearly germane to it.

There is an extensive literature on testimony, much of it a debate between 'reductionists',¹⁸ who do not see testimony as an autonomous source of epistemic authority, and 'anti-reductionists',¹⁹ who typically see testimony as a source of warrant in itself, not reducible to, say, perception or memory (Adler, 2017).

Previously²⁰ we argued that if beliefs in the continued uniformity of nature or the existence of other minds qualify as foundational then our evidence principle

¹⁷ From p23.

¹⁸ See for example Hume (1977 [1748]:74-5); Hume (1978 [1740]: Bk I part III section IV); Adler (2002:136-159); Fricker (1987); Fricker (1994); Fricker (1995); Fricker (2004); Fricker (2006).

¹⁹ See for example Reid (1983:94-95); Audi (1997); Coady (1992); Burge (1993); Burge (1997).

²⁰ p88ff.

cannot disallow them if doing so would undermine the social context on which moral life, and therefore our evidence principle, depend. A similar ('anti-reductionist') argument could apply to testimony in principle: social life would be unthinkable without shared social projects, which would be impossible if no testimony-based beliefs were allowed. Not believing anything another person tells me is tantamount to not trusting him; not believing anything anyone tells me means trusting no one. Shared social life and shared social projects rely not only on shared knowledge and belief, but on a threshold of trust between community members, extending to institutions and institutional artefacts like law and science.

Not all testimony is trustworthy of course. The issue is knowing whom and when to trust and not trust. EP4 cannot force us to discount all testimony on principle if this would undermine the social and moral context which EP4 presupposes. But nor can we assume all testimony is true: testimonies sometimes contradict each other; and people can lie, cheat, mislead or be innocently mistaken.

From the 'reductionist' side Adler argues that we trust testimony not only because we have no alternative but because 'acceptance of testimony has strong empirical support' from our 'background beliefs' (2002:136). Just because a 'belief *originates* in testimony' does not mean it is forever '*sustained* by it' (2002:137), or at least not by its originating testimony. I may first hear of a serious crime from my neighbour, but '[a]lmost immediately there is verification from many other sources, in particular, newspaper and TV reports' (2002:138). While each may involve its own testimonial chain, their 'convergence corroborates ... the report's accuracy, as well as my reasons for accepting it'.

Adler questions not whether trust is involved or needed but whether, 'where trust is needed, it is to compensate for a lack of knowledge (justification, or adequate reasons)' (2002:139). He aims to 'diminish the gap between our dependence on testimony and the fragile nature of testimonial trust' (2002:141). His chosen context is the 'null setting', where

the hearer has no specific information concerning his informant's reliability or trustworthiness[,] ... [t]he costs of error from false testimony are not very

high, and seeking or acquiring the information offered through testimony in other ways is costly. Otherwise testimony loses its driving force as a device of economy (2002:142).

He suggests three possible positions on whether it is correct to accept an informant's testimony in this context:

1. The neutral position: that trust or acceptance should be judged on a case by case basis. There is no presumption favoring or disfavoring acceptance of testimony. [...]

2. The negative-bias position: that one should always regard the word of another critically. It is suspect, until proved otherwise[.]

3. The positive-bias (or default) position: that one ought simply to accept a speaker's testimony unless one has special reason against doing so.

Adler thinks the positive-bias, default, position

fits our testimonial practices ... since it is not credible that we would remain so heavily dependent on testimony were it not able to secure for us overwhelmingly useful (true, relevant) information (2002:143).

This position is a 'practical necessity' (2002:142) but also has 'empirical backing' (2002:144). First is the 'enormous evidence, from our earliest years, of reliable testimony in its basic role of conveying information' which we verify 'simply by acting on it' (2002:148). Second are the 'powerful institutional and social constraints on us to speak truthfully and reliably—most prominently, reputation' (2002:149). Third is our 'knowledge of teaching and learning in regard to testimony in our community' (2002:150). Fourth is our 'background knowledge' which grounds our

judgments of prior plausibility because testimony, to be minimally acceptable, must not be dissonant with the hearer's beliefs (2002:151).

Lastly, we use testimony ‘largely under conditions that risk detection of error (or falsification)’ (2002:153). Our default positive-bias rule ‘reflects a resilient history of overwhelmingly reliable testimony’.

The default rule’s ‘hold’ (2002:154) on us is not just evidential. Its ‘pull’ also derives from ‘social norms of civility and courtesy’. But ethical and evidential pull together:

The practical economies secured by the default rule are respected because we have good reason to expect that the benefits are not purchased at the cost of incurring serious error. The central pillar of that rationale can reside only in our background knowledge. (2002:155)

Adler concludes that our background beliefs provide ‘enormous empirical support for the acceptance of testimony’ (2002:159): except of course when we have good reason to override the default, which the positive-bias position allows for.

Uniformity of nature

Adler’s rationale is based on an inductive argument, which in turn presupposes belief that nature was uniform in the past and will continue uniform in the future. This is part of Clifford’s justification for accepting Thucydides’ account of the siege of Syracuse:

Our experience is that manuscripts exist which are said to be and which call themselves manuscripts of the history of Thucydides; ... We find also that men do not, as a rule, forge books and histories without a special motive; we assume that in this respect men in the past were like men in the present; and we observe that in this case no special motive was present. (1877, 1879:208-9)

It might seem counter-intuitive that testimony-based belief about the past presupposes belief in the future uniformity of nature. But we project forward by

necessity, because although 'life must be understood backward ... it must be lived forward' (Kierkegaard, 2000:12). The positive-bias rule is applied to each new instance. The potential believer applies the rule, implicitly or explicitly, to each new case to decide whether to accept the testimony or not, and each time presupposes that nature will continue to be uniform. Otherwise we could only say we *were* right to accept what Thucydides wrote, not that we are right and will continue to be. We must at least assume all current evidence justifying our confidence in Thucydides' account will continue to exist into the future.

But if that is true of a belief about the past based on testimony, something similar applies to a belief based on personal memory. Most of the time we simply assume our memories are accurate. But we know they are not error-proof because people can remember the same event in different and incompatible ways. Despite this we implicitly or explicitly apply a default positive-bias rule to accept the truth of our personal memories. With rare exceptions we are rewarded by overwhelming confirmation that this rule is sound, justifying our continued application of it to each new memory-based belief. So even with memory-based beliefs about the past we assume nature will continue to be uniform, and that this overwhelming confirmation will continue.

If every belief about the past is based on personal memory or testimony or (most likely) a combination of the two, then it seems a belief about the past is always implicitly a belief about the future, to the extent that a belief about the past contains within itself an assumption that the future will resemble the past. We will now turn to beliefs about the future themselves.

Beliefs about the future

Kate says her train will arrive at 5 pm, so Adler goes to meet her at 5 (2002:8). He asks if he knows Kate will be there at 5, but we could equally ask if he has sufficient evidence to believe this. In a sense he cannot know: there could be a rail strike; a sudden snowfall; she could miss her train; she or the station could auto-destruct. Does that mean he cannot have sufficient evidence that she will be there? Applying EP4 he would then be morally wrong to believe she will be there.

But would he have sufficient evidence that any of these obstacles is likely enough to make him believe Kate would *not* be at the station at 5? Are the railways going through a bad patch? Snow or industrial action forecast? Does Kate have a history of unreliability? If not then his belief that Kate will be at the station at 5 is surely justified, because the contrary belief that Kate will not be at the station at 5 is unjustified?

A third option is neither believe nor disbelieve. If there is insufficient evidence either way, would the (morally) right thing be to remain agnostic?

We could argue that agnosticism only makes sense from a purist perspective. 5 pm arrives, Kate is there and you meet her, despite neither believing nor disbelieving she would be there. Are you surprised by this 'coincidence'? (Adler, 2002:8) What if Kate had previously been unreliable, if snow was forecast, with threats of industrial unrest; and you still turned up at 5 – because you said you would, again neither believing nor disbelieving she would be there – would you be surprised if she was there or if she was not?

We must separate a few strands. This is a belief about a future event which has not happened, but which in many significant respects is like events which have happened before. To that extent we could substitute an inorganic event like a solar eclipse. But it is also an event involving humans with free will; and possibly implicit or explicit promises and/or consciously shared expectations.

The involvement of free will does not necessarily invalidate inductive inference. Adler feels justified in believing the mailman will greet him even though the mailman is free not to:

Free will ... excludes neither regularities governing nor ... knowledge of human behavior (2002:140).

This applies to Kate's reliability and to collective human action on which we can base beliefs about predictability in rail services and the likelihood of strikes. Some

at least of the contingencies impacting Kate's punctuality could themselves be inferred, in principle, from prior physical states and scientific laws presupposing the continued uniformity of nature: mechanical characteristics of rolling stock, seasonal weather patterns and so on.

Earlier we offered reasons for considering belief in the continued uniformity of nature to be justified, either as foundational²¹ or as a presupposition²² for applying tacitly confirmed background beliefs. We also argued for adopting a default positive-bias rule for testimony in general,²³ to include relevant aggregate knowledge relating to the physical systems which Adler hopes will bring Kate to the station by 5. If on balance he has no good reason to expect negative contingencies, then – for now excluding free will – he would seem justified in believing Kate will be at the station at 5. This is because the consequences of believing otherwise, including remaining agnostic, should be considered as if universalised across Adler's entire community. Shared projects and social life in general would be impossible if no one felt justified in holding beliefs about the future based on inductive inference from collective past experience relating to the physical world.

When we readmit free will we get a similar result, but for slightly different reasons. There is the aggregate effort by railway staff following established systems and procedures, plus Kate's reliability record. There is no need to assume perfection, just past performance levels high enough to discount any positive reason for thinking Kate will not arrive by 5. To assess whether Adler would be justified in remaining agnostic we should again extrapolate the same policy across the whole community.

In this thought experiment a putative justification for agnosticism is that Kate might possibly not be at the station at 5, but not because of unreliability on her part or an inefficient railway service. Nothing either Kate or any railway employee can do could remove the sceptic's doubt. So if his doubt is justified so would anyone

²¹ p88.

²² p90.

²³ p94.

else's about any other planned outcome involving an element of free will. No person or group could do anything to create an effective reputation for reliability, because agnosticism would ignore track record. No matter how mild, peaceful, considerate and loving X had behaved in the past his acquaintances could still have insufficient reason for discounting the possibility that he, and indeed anyone else, could turn psychopathic the next instant.

From EP4's perspective we would argue there would be at least a 'conflicting and overriding moral imperative' not to undermine social life by insisting on agnosticism in principle, if not that the possibility of having legitimate expectations about people's future behaviour would be a pre-requisite for moral life, and therefore for the EP4 condition 'if anything is morally [right or] wrong'.²⁴

We will now see how Clifford treats inductive inference on the assumption of the uniformity of nature, as it leads to insights not just on social and moral life but on belief itself.

Belief as a guide to action

A critic could claim that believing in the continued uniformity of nature is to 'believe ... upon insufficient evidence' and therefore breaches CP. We might have 'sufficient evidence' that nature *has* been uniform. But to infer from this that nature will always be uniform would presuppose what we want to infer: that nature is and will be uniform. The classic source for this 'problem of induction' is Hume (1969: Book I, Part III, section VI); but Ayer offers a convenient summary:

In [inductive] reasoning we make the assumption that there is a measure of uniformity in nature; ... that the future will, in the appropriate respects, resemble the past. ... But, as Hume pointed out, this assumption is not demonstrable; the denial that nature is uniform ... is not self-contradictory. Neither, as Hume also saw, is there any means of showing, without logical circularity, that the assumption is even probable. For the only way of

²⁴ This particular belief might also qualify as a self-fulfilling belief, discussed in the next section: p110ff.

showing that it was probable would be to produce evidence which confirmed it, and it is only if there are fair samples in nature that any evidence can be confirmatory. But whether there are fair samples in nature is just the point at issue. (1956:72)

Clifford claims however that it is 'in the very nature of belief' to go 'beyond experience', so the question is not 'May we believe what goes beyond experience?' but 'How far and in what manner may we add to our experience in forming our beliefs?' (1877, 1879:206). He offers not an inference but a 'rule' of 'extreme simplicity and vast practical importance' (1877, 1879:205):

We may believe what goes beyond our experience, only when it is inferred from that experience by the assumption that what we do not know is like what we know. (1877, 1879:210)

For example spectroscopic evidence supports the belief that the sun contains hydrogen. This belief cannot be accepted as true, so cannot 'help in the right guidance of human action' if it is 'accepted on unworthy grounds, and without some understanding of the process by which it is got at' (1877, 1879:208).

But when we understand this process as the 'ground of the belief' it becomes a 'very serious and practical matter'. If the sun contains no hydrogen then the spectroscope is an 'uncertain guide' for chemical analysis – so for consistency we should no longer trust spectroscopy generally, which up to now has 'enriched us not only with new metals ... but with new processes of investigation, which is vastly greater' in value (1877, 1879:208).

The nature of belief itself involves the 'possibility of inference' (1877, 1879:210), and so the assumption of uniformity in nature is implicit in all non-tautological beliefs that go beyond experience:

No evidence ... can justify us in believing the truth of a statement which is contrary to, or outside of, the uniformity of nature. If our experience is such that it cannot be filled up consistently with uniformity, all we have a right to

conclude is that there is something wrong somewhere; but the possibility of inference is taken away; we must rest in our experience, and not go beyond it at all. (1877, 1879:210)

Inference would be impossible because inference relies on assuming nature is uniform. An event which was really outside the uniformity of nature would be one only 'those whose actual experience it was' would have the right to believe; and 'no inference worthy of belief could be founded upon it' (1877, 1879:210).

Although we can envisage someone witnessing an event which appeared to contradict the uniformity of nature, no one else could infer a legitimate belief from it (for example that they would also experience the same thing under the same circumstances), because the right to infer would presuppose the uniformity of nature in every other necessary respect.²⁵

Despite this though we are not 'bound to believe' – indeed we 'have no right to believe' – that 'nature is absolutely and universally uniform':

*The rule only tells us that in forming beliefs which go beyond our experience, we may make the assumption that nature is practically uniform **so far as we are concerned. Within the range of human action and verification, we may form, by help of this assumption, actual beliefs; beyond it, only those hypotheses which serve for the more accurate asking of questions.***

[Emphases added.] (1877, 1879:210)

The justification of this rule, including both its extent and its limitations (emphasised above) derives from the role belief plays in the 'right guidance of human action':

...if [a belief] is to be used as a guide to action, as a hint of what the future is to be, it must assume something about that future, namely, that it will be

²⁵ For example that people generally tell the truth and generally know when they are and are not doing so; that people's memories generally persist reasonably accurately and when they do not we know this by comparison with other records and traces which persist unchanged; that deliberate physical records of events and contingent traces of their impact (ashes from fires, debris from explosions, puddles after showers etc) persist in known and reliable ways; and so on.

consistent with the supposition that [the event in question] really took place yesterday; which is going beyond experience. Even the fundamental 'I am,' which cannot be doubted, is no guide to action until it takes to itself 'I shall be,' which goes beyond experience. (1877, 1879:206)

The assumption is not just that 'the unknown fire of to-day is like the known fire of yesterday', but that the future will behave such that the remembered and recorded past will persist as a single and unified past with continued connection to present and future. Believing 'what goes beyond experience' is 'involved in the very nature of belief', leaving no alternative but to 'add to our experience on the assumption of a uniformity in nature' (1877, 1879:206).²⁶ Belief is thus intrinsically forward-looking – which would extend our working definition of 'belief'²⁷ as 'what we regard as true'. Clifford insists a belief has potential 'influence upon the actions of him who holds it' (1877, 1879:181). Peirce agrees, claiming the 'essence of belief' is the 'establishment of a habit' (1878, 1998:144) – 'habit' being short for 'rule of action' (1878, 1998:143) – which must presuppose there will be a future resembling the past.

Belief as intrinsically forward-looking

We should unpack the claim that belief is intrinsically forward-looking. Could someone not have a belief only about the past? The question seems to be whether we could separate a belief from anything which might be done with it: because whatever will be done with it will by definition be in the future. This in turn seems related to whether we can separate the belief from the believer whose belief it is. Could we for example see a belief as a 'trace' or record of an event, like a setting within a mechanical memory (or a snapshot, or an edge of wet sand after the tide has turned) but without any implication of any future event triggered by it? A true belief at least seems to share some features of this kind of trace. But a belief can also be false, and if only to reflect this possibility it is hard to see how we can remove the believer from the picture. For this trace to be something which

²⁶ Also: *...it is our duty to guide our beliefs by inference from experience on the assumption of uniformity of nature and consciousness in other men, and by this only. Only upon this moral basis can the foundations of the empirical method be justified. (Clifford, 1875, 1879:175)*

²⁷ See p22.

could be either true or false it seems it must count as potential information, or at least something which could be put to use, by someone or something capable of putting it to use, even if just to infer some other belief or proposition.

For a less abstract approach consider my belief that I walked upstairs five minutes ago. A belief is something I take to be true, so I take it to be true that I walked upstairs five minutes ago. Is part of taking this to be true that if someone asked me a minute from now whether I walked upstairs six minutes ago I would say yes? So taking it as true that I came upstairs five minutes ago implies I will continue to take it as true, so after a minute I will believe I walked upstairs six minutes ago?

If yes, is this a characteristic of my belief or of myself as a believer? If I did not project my belief forward into the future would that be because it was not a genuine belief, or because it was not a genuine belief for me, or because I would not be behaving as a genuine or 'normal' believer? And if it was a characteristic of me as a believer, is this because, as the existentialists argue, I am 'a being whose existence comes before its essence'; someone who is 'nothing other than what he makes of himself'; something which 'projects itself into a future, and is conscious of doing so' (Sartre, 2007:22-3); because as 'an existing individual' I am 'constantly in the process of becoming' (Kierkegaard, 1974:79)?

It is hard to disentangle these questions, because it is hard to envisage a belief without a believer; or to envisage a believer except as someone whose life 'must be lived forward'. So one question is whether it is possible to separate the belief from the believer. In the context of EP4 however we do not need to separate the two, as EP4 is a moral principle aimed at a believer's behaviour in respect of his beliefs.

A related question though is whether a belief can be separated from a believer's other beliefs. We were trying to restrict ourselves to a belief with no future commitment. But we can ask if the belief that I came upstairs five minutes ago can really stand as an atomic belief independent of any related beliefs at least some of which do seem to have a forward-looking aspect. There are beliefs about what 'upstairs' means in terms of hypothetical futures: that if I were to look out of the

window I would see I am higher up than when I looked out of the window downstairs, and that to go to the kitchen I will have to go back downstairs. A reason for thinking my 'atomic belief' cannot stand independent is that perhaps I only understand its propositional content because I have other beliefs about (say) the configuration of my house, at least some of which have implications for the future: that the house will stand secure and not collapse under its own weight; that the stairs are strong enough to bear my weight and the timber will not liquefy. Otherwise I would not have ventured upstairs: I may not have entered the house because it would have been unsafe. Included in my belief that I walked upstairs five minutes ago may be that I did it knowing I was safe and the staircase would not collapse and leave me stranded. If, when I believed I came upstairs five minutes ago, I was truly agnostic about the future awaiting me, then my belief would have been like a belief in a dream, when anything could happen next, and I could suddenly find myself on a mountaintop. Again it is hard to separate characteristics of my belief about coming upstairs five minutes ago (plus any related beliefs) from characteristics of myself as a believing agent. Certainly my conception of what a belief is does not seem independent of its being held by a believing agent, and again in the context of EP4 we can see the two as inseparable.

Another way to interpret the intrinsically forward-looking nature of belief could be that at least some of our beliefs must be forward-looking, so it is intrinsic to belief that a belief can be forward-looking. For example we could ask: what would our beliefs be if we did not use at least some of them as guides to (future) action? If we still engaged in intentional behaviour but did not use our beliefs as guides, how would we translate our intentions into actions which carried out those intentions? How would we achieve anything? And if we must use at least some of our beliefs as guides to action then is it possible to have a belief which could not be used as a guide to action? If not, that would suggest that all our beliefs must be available for use as guides to action. And if something *can* be a guide to action, then surely it *is* a guide to action?

The idea of belief as a guide to action (or at least enough of this idea as our current context needs) may actually be implicit in the idea of belief as 'what we

regard as true'. It would be hard to understand how we can regard something as true without that 'regard' having persistence and practical consequence. Three ostensibly separate ideas seem interconnected, if not aspects of the same thing:

(3.4.1) Belief as 'what we regard as true'.

(3.4.2) Belief as a guide to action.

(3.4.3) Belief as presupposing continuity and shared features between past, present and future.

To regard it as true (3.4.1) that I have one dog and he is white, that belief must be identifiable against a universe of alternative and potential beliefs. I must believe I know what a dog is, and what it is for a dog to be the colour he is and not another colour. I would have some idea about when a white dog entered my life, as this would be part of what it is to believe he is my dog. If I regard it as true that I have one dog and he is white then I would not recognise a black dog as my dog, as I would not regard it as true that I also have a black dog. Countless other dispositions, implications and possible scenarios are reflected in the idea of belief as a guide to action (3.4.2). For present purposes that is all (3.4.2) needs to contain. My claim to regard it as true that I had one dog and he was white would be viewed with suspicion if I seemed not to know what a dog was, as something that could be a particular colour; or if I seemed not to know when a white dog entered my life. These would be hard to articulate except by presupposing continuity and shared features between past, present and future (3.4.3).

A critic could complain we have deliberately chosen a belief which only makes sense if (3.4.1–3) all apply together. But the challenge would be to find an example where (3.4.1) applies without (3.4.2) or (3.4.3). It would need enough of the relevant logical, phenomenological and epistemological features for (3.4.1) to apply, but without the practical and phenomenological features which (3.4.2) and (3.4.3) need. If the idea of belief as 'what we regard as true' (3.4.1) and as incorporating some sense of 'guide to action' (3.4.2) presupposes continuity and resemblance between past, present and future (3.4.3), then it seems we can

assume that 'nature is practically uniform so far as we are concerned' (Clifford, 1877, 1879:210).

The proposition that nature is, was and will be uniform must have the status of a rule or presupposition for an 'evidentialist' principle like CP because, as we will see, if it counts as a belief among other beliefs then this kind of evidentialism might lead to scepticism. Although we have argued CP is implicitly a moral principle, Clifford worded it in absolute, universal terms without caveats or qualifications. And while he gives a moral justification for CP in terms of the social consequences of obeying and disobeying it, he does not capitalise on the social preconditions for morality per se. The justification he claims for treating the uniformity of nature as a rule or presupposition rather than a belief is not on moral grounds but because it is implicit in the nature of belief. What he could have done, but did not do, was make this explicit in CP: '*Assuming (contingent) belief to be possible at all, then it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence*'.

But even this manoeuvre might not succeed. Whether the assumption of the uniformity of nature is implicit in the nature of belief or CP had an explicit assumption about the possibility of (contingent) belief, neither would necessarily legitimise the assumption of the uniformity of nature. This is because an alternative conclusion could be that if we adopt an evidentialist principle like CP then we should not believe anything except perhaps necessary truths.

An 'anti-evidentialist' could for example claim that, if evidentialism is true, the idea of belief as a guide to action is illegitimate. If evidentialism presupposes that belief is intrinsically forward-looking and that there will be a future resembling the past, then so much the worse for evidentialism. Not that we cannot justify assuming the uniformity of nature without circularity, but that *evidentialism* requires us to justify assuming the uniformity of nature without circularity, and on the basis of evidence. For example:

- (3.5.1) According to evidentialism it is wrong to believe anything on insufficient evidence.

- (3.5.2) A belief about the future or the past presupposes the future will resemble the past.
- (3.5.3) There is insufficient evidence that the future will resemble the past.
- (3.5.4) Therefore [from (3.5.1) and (3.5.3)] according to evidentialism it is wrong to believe the future will resemble the past.
- (3.5.5) Therefore [from (3.5.2) and (3.5.4)] according to evidentialism it is wrong to hold any belief about the future or the past.
- (3.5.6) Therefore [from (3.5.1) to (3.5.5)] evidentialism leads to scepticism about the future and the past.

We could continue:

- (3.5.7) The idea of belief includes that of a guide to action.
- (3.5.8) The idea of belief as a guide to action presupposes there will be a future which will resemble the past.
- (3.5.9) But [from (3.5.5)] according to evidentialism it is wrong to hold any belief about the future or the past.
- (3.5.10) Therefore [from (3.5.8) and (3.5.9)] according to evidentialism it is wrong to hold any belief whatsoever.

Clifford would be unlikely to accept this argument, although he would accept at least some of its premises. He would presumably accept (3.5.1) and – with a key reservation – possibly both (3.5.2) and (3.5.3). The reservation would either stop him accepting (3.5.4) or allow him to accept only a qualified version of it.

The reservation is whether ‘the future’ and ‘the past’ in (3.5.2–4) refer respectively

to everything the future could bring across all reality and everything which has ever happened or existed across all reality, or only refer to aspects of reality relevant to the specific belief in question. He seems to draw this distinction when he claims we 'have no right to believe' that 'nature is absolutely and universally uniform' (1877, 1879:210) but that 'we may make the assumption that nature is practically uniform as far as we are concerned' and '[w]ithin the range of human action and verification, we may form, by help of this assumption, actual beliefs'.

But Clifford does not give a justification for thinking (contingent) belief must be possible. His justification for CP is largely based on the moral impact of obeying or disobeying it in a social context. So he could have suggested a similar rationale to justify adopting the 'rule' that nature will continue to be uniform: because of the social and moral consequences of not adopting it.

In the case of EP4 however we make this explicit. EP4 is expressed in moral, not epistemic terms. If the belief that the future will resemble the past qualifies as a descriptive belief, then it would be morally wrong to believe, without sufficient evidence, that the future will resemble the past. But the idea of deliberate action,²⁸ let alone social and moral life, presupposes the future will resemble the past. So if we say it is morally wrong to form beliefs on the assumption that the future will resemble the past because it is morally wrong to believe (on insufficient evidence) that the future will resemble the past, then we are saying it is morally wrong to hold any belief which presupposes there will be a future which resembles what we understand and remember as the past. It would be morally wrong to believe in the existence of other people and other minds with intentions, expectations, hopes and fears; morally wrong to believe we have any future, or any past which will continue to make sense to us in the future; morally wrong to believe anything the possibility of moral life depends on.

But EP4 explicitly presupposes the moral domain is possible and exists: '{If anything is morally wrong, then...}'. The qualification '[except when the unjustified believing is outside the believer's voluntary control]' would admit an involuntary

²⁸ Clifford's "fundamental 'I am'" (1877, 1879:206).

over-belief that the future will be like the past. And '{...in the absence of any conflicting and overriding moral imperative}' would admit a voluntary over-belief that the future will be like the past, if that is a necessary condition for morality – which it seems to be.

Conclusion

We argue that expressing EP4 in explicitly moral terms (which Clifford left implicit) insulates it from radical-sceptic challenges like (3.5.1–10). Even if successful from a purely epistemic perspective, the sceptical alternative would reduce the domain of 'justified belief' to a vanishing point coextensive with the sceptic's own domain of 'certain knowledge'.²⁹ A belief could not be a guide to action, because we could not legitimately predict the future or recollect the past. A belief could still be what we regard as true, but only things we know for certain would qualify: the Cartesian cogito, analytic truths etc – assuming these propositions still made sense in such a whittled-down world. But even if radical scepticism is impregnable as an epistemological stance, it can be defeated on moral grounds as a life option or orientation to the world, because if carried through consistently it would outlaw belief structures on which social and moral life depend. This is enough to save EP4.

Our opponent might question why an anti-sceptical evidence principle allows us to dismiss agnosticism in the context of beliefs about the future but embrace it when it comes to (say) belief in God. **Chapter 4** will discuss religious belief in general. For now we can say that with beliefs about the future the past evidence exists and future evidence will eventually exist. The issue is whether we are justified in assuming uniformity in nature. If not then all our beliefs might be doomed, certainly all beliefs containing an element of projection into the future, including those on which social and moral life depend. This has no parallel in respect of belief in God, even if our concept of God includes the belief that God is creator, maintainer and guarantor of reality. Our belief in the existence of reality is not dependent on the belief that God is its creator, maintainer and guarantor. That latter belief would be just another descriptive belief needing evidence. It would not

²⁹ Picked up again in **Chapter 5**, especially p261ff.

be a precondition of (future-directed, contingent) belief per se, which assuming uniformity of nature is. Nor is it a precondition of social or moral life, as many stable communities function without belief in God. For now we conclude that the assumption of the uniformity of nature, and therefore any suitably qualified beliefs about the future and the past, would correctly survive EP4. **Chapter 4** will address whether, and under what circumstances, belief in God would survive EP4.

3.4 Personal relations and self-fulfilling beliefs

Faith in a fact

The previous section included Adler's story about meeting Kate at the station.³⁰ It did not consider whether Adler's belief in Kate's reliability influenced her behaviour in any way, which could make it another of James's potential counter-examples to CP: 'questions of fact ... concerning personal relations, states of mind between one man and another' (1896, 2000:213), where 'faith in a fact can help create the fact' (1896, 2000:214):

Do you like me or not? ... Whether you do or not depends ... on whether I meet you half-way, am willing to assume that you must like me, and show you trust and expectation. The previous faith on my part in your liking's existence is in such cases what makes your liking come. But if I stand aloof, and refuse to budge an inch until I have objective evidence, ... ten to one your liking never comes. (1896, 2000:213)

Two separate but overlapping categories: beliefs specific to interpersonal relations, and 'self-fulfilling' beliefs. The belief that someone you have not yet met will like you could belong to both. Self-fulfilling beliefs may occur in the context of an interpersonal relationship or may involve only one person:

Suppose ... you are climbing a mountain and have worked yourself into a position from which the only escape is by a terrible leap. Have faith that you can successfully make it, and your feet are nerved to its accomplishment.

³⁰ p96ff.

But mistrust yourself, and think of all the sweet things you have heard the scientists say of maybes, and you will hesitate so long that, at last, all unstrung and trembling, and launching yourself in a moment of despair, you roll into the abyss. (James, 1897, 1912a:59)

Are self-fulfilling beliefs another legitimate exclusion, like prescriptive beliefs?

Closer examination suggests they could support rather than undermine EP4. James's 'questions of fact' are beliefs about the future. In the first case, I would seem to have no evidence whether you will like me or not, because we have never met. But that is no reason for not acting as if you will like me, particularly if I have reason to think you are more likely to if I act as if I think you will. Observed human behaviour generally supports that assumption. So I might be correctly believing, on *sufficient* evidence, that you are more likely to like me if I act as if I think you will:

*The desire for a certain kind of truth here brings about that special truth's existence; and **so it is in innumerable cases of other sorts.** [Emphasis added.] (James, 1896, 2000:213)*

'Faith', 'trust', 'expectation' and 'desire' are neither synonyms of nor inseparable from belief – although separating them may require 'engag[ing] in some abstraction from the natural workings of belief' which apply in social and conversational contexts where 'truthfulness' may be a lower aim than 'tactfulness' (Adler, 2002:14) – or indeed optimistic faith or generosity of spirit.

Compare the scenario where Kate and the railway service are known to be reliable and good weather is forecast.³¹ Because EP4 presupposes a future resembling the past, we concluded Adler would be morally right to believe Kate would be at the station at 5 pm, and that any other option would be morally wrong if it meant ignoring the evidence. In this case we are assuming his belief did not influence Kate's punctuality.

³¹ p96ff.

But believing someone I have never met will like me could influence the outcome. If I believe it is more likely than not that the person will like me, and I behave in a friendly and engaging way which makes it even more likely, then I am believing on the basis of the available evidence – ‘on the assumption of a uniformity in the characters of men’ (Clifford, 1877, 1879:209). James says Clifford would ‘stand aloof, and refuse to budge an inch’ until he has ‘objective evidence’ (1896, 2000:213), sooner dying on the mountaintop than make a leap of faith:

[S]uppose that, having just read the Ethics of Belief, I feel it would be sinful to act upon an assumption unverified by previous experience, why, then I shall hesitate so long that at last, exhausted and trembling, and launching myself in a moment of despair, I miss my foothold ... (James, 1897, 1912b:96-7)

But Clifford explicitly acknowledges those

many cases in which it is our duty to act upon probabilities, although the evidence is not such as to justify present belief; because it is precisely by such action, ...that evidence is got which may justify future belief. So that we have no reason to fear lest a habit of conscientious inquiry should paralyse the actions of our daily life. (1877, 1879:189)

This may not be exactly how James would describe the same adventure, but there seems little difference in principle between James’s ‘hesitate so long’ and ‘despair’ and Clifford’s ‘paralyse the actions of our daily life’.

We will now change the first story. What if most people I meet for the first time end up *not* liking me? Should we then argue that, other things being equal, I would be morally wrong to believe the next new person I meet would like me because I would be believing against the evidence? But other things may not be equal: it might be morally better to trust both of us to defy the odds to make an eventual good relationship more likely.

It is unsurprising that interpersonal relations are a realm where moral imperatives of (say) trust and generosity can outweigh those concerning over-belief. Conflicting moral imperatives are not counter-examples to EP4. Belief and trust are different, but in the quintessentially moral domain of interpersonal relations an instance of believing could also be an instance of trusting. (Lying to the Gestapo may also be an instance of courage and of loyalty.³²)

All these cases could qualify as 'positive thinking' – where, in either an interpersonal or solitary context, what might otherwise count as over-belief encourages a desirable outcome. We are currently assuming such beliefs are possible, which we will evaluate later.³³ But it is worth revisiting Adler's story in **Chapter 1** about Jim,³⁴ who 'cannot take himself to believe that he is handsome as a way to avoid depression' (however 'rational') because 'the lessened depression does not bear on the truth of whether he is handsome' (2002:10). As 'handsomeness' can be in the eye of the beholder we will substitute a purely descriptive belief:

(3.6) It would be contradictory for me to think that (i) I believe I have £1 million in the bank; (ii) my only reason to believe I have £1 million in the bank is that it will make me happy; and (iii) my happiness has no bearing on the truth of whether I have £1 million in the bank.

For Adler, my condition of 'full awareness' makes me conceptually unable to believe (i) if I believe that (ii) and (iii) are also true, which (3.6) assumes they would be. I would be misunderstanding what 'believing' was: perhaps confusing it with 'wishing'.

Contrast:

(3.6.1) It would *not* be contradictory for me to think that (i) I believe I have £1 million in the bank; (ii) my only reason to believe I have £1 million in

³² p22.

³³ p115.

³⁴ See p26.

the bank is that my bank statement has a credit balance of £1 million;
and (iii) the £1 million credit balance on my bank statement *does* have
a bearing on the truth of whether I have £1 million in the bank.

In (3.6.1) the credit balance ((ii), (iii)) is evidence for (i) my belief about having £1 million in the bank. In (3.6), my happiness is not evidence for the same belief.

We can now apply this to ‘Do you like me or not?’:

(3.6.2) Would it be contradictory for me to think that (i) I believe you will like me; (ii) my only reason to believe you will like me is that it will tend to make you like me; and (iii) this eventual tendency for you to like me has a bearing on the truth of whether you will like me?

This is less cut-and-dried, hence the question format. It seems coherent for me to believe you will like me, and therefore for me to think I believe you will like me. It also seems coherent to have, as my only reason for believing you will like me, that this will tend to make you like me.³⁵ Could your eventual tendency to like me have a bearing on the truth of whether you will like me? Surely yes, although perhaps not in Adler’s original sense. Does your eventual tendency to like me qualify as *evidence* for my belief that you will like me? That is questionable. It would certainly count as evidence supporting a belief like:

(3.6.3) I believe that if I believe (if, that is, I succeed in making myself believe) you will like me, then this will tend to make you like me.

(3.6.3) seems a paradigm case of ‘faith in a fact’. The word though is ‘faith’ not ‘belief’. Our scope does not warrant a detailed analysis of the difference between the two, beyond acknowledging that ‘faith’ seems intuitively appropriate for beliefs like these which, successful or otherwise – and assuming they are possible – are intended as self-fulfilling.

³⁵ Scott Aikin calls these ‘doxastically efficacious’ beliefs (2014:85).

Applying EP4 to these cases yields ostensibly sound results. In (3.6) it could be morally wrong to believe without evidence that I have £1 million in the bank: I could act on that belief and impoverish my family by spending money I do not have. The equivalent belief in (3.6.1) passes EP4 because here I do have sufficient evidence. Belief (i) in (3.6.2) that you will like me arguably also passes EP4. In theory we could consider belief (i) outside (3.6.2), when it might seem to fail EP4. But moral principles like EP4 apply to acts and omissions in their context. So we cannot remove the belief from its context without falsifying what we are using EP4 to evaluate. Key to that context is my reason – (ii) in (3.6.2), and spelt out in (3.6.3).

We are seeing (3.6.3) as a belief for which we could have sufficient evidence, generalising from observations of human nature. When discussing involuntary beliefs³⁶ we stressed that our principle (now EP4) applies to aspects of doxastic behaviour under voluntary control, even if we have only indirect control over the beliefs themselves. We will therefore use ‘voluntary belief’ as a convenient label for anything we can voluntarily do to increase the likelihood of having a belief we might want to have, including behaving as if we hold that belief. This should shield us from objections like Williams’s ‘if I could acquire a belief at will, I could acquire it whether it was true or not’ (1973:148). Since (3.6.3) would be a belief supported by sufficient evidence, my ‘voluntary belief’ that you will like me – (i) in (3.6.2) – would be permitted by EP4: because it would not be believing ‘on insufficient evidence’. And even if my experience is that most people I meet for the first time end up *not* liking me, EP4 would still permit my ‘voluntary belief’ that you will like me if there is an overriding moral imperative to make an eventual good relationship more likely.

Our conclusion is that so far ‘faith in a fact’ cases, in both individual and interpersonal contexts, seem unproblematic for EP4. **Chapter 4** will look at James’s apparent attempt to extend ‘faith in a fact’ to religious belief.³⁷ For now we will briefly consider a related category of beliefs arising in interpersonal contexts but without qualifying as ‘self-fulfilling’ because although there is an

³⁶ p76.

³⁷ See p165.

element of intention the intended outcome is not exactly the truth of the belief.

Meiland (1980) describes scenarios which he thinks violate the 'normative principle that one should believe whatever is backed by sufficient evidence'.³⁸ We will cover the first and introduce the second here, but will need to postpone further discussion until **Chapter 5**.³⁹

Smith and Jones

Smith and Jones have been 'business partners and exceptionally close friends' for many years. Jones now discovers 'evidence which is sufficient (in anyone's eyes) to justify the belief' that Smith has been embezzling. But to protect both partnership and friendship Jones 'decides not to believe that Smith stole money from the firm. In fact ... he decides that Smith did not steal money from the firm' (1980:15). This would not be a 'self-fulfilling' belief because Jones is not believing in Smith's innocence to make it true that Smith is innocent, however much Jones might want that to be true. In a non-magical world that would be impossible anyway, despite positive belief helping you leap a chasm. Meiland tells the story in terms of 'positive obligation' rather than negative obligation or positive permission.⁴⁰ But we can recast it in line with EP4:

- (3.7) {If anything is morally wrong, then} it is [morally] wrong of Jones to believe [knowingly or irresponsibly] that Smith is innocent of embezzlement on insufficient evidence {in the absence of any conflicting and overriding moral imperative} [except if the unjustified believing is outside Jones's voluntary control].

We are taking it that if Jones has sufficient evidence of Jones's guilt, then he has insufficient evidence of his innocence.

This scenario should not undermine EP4. Jones would not be morally wrong to

³⁸ Meiland's scenarios may also qualify for treatment as 'moral encroachment' cases, which are outside our scope. See p295.

³⁹ p246ff.

⁴⁰ See p13.

believe Smith is innocent if there is a 'conflicting and overriding moral imperative' to protect the business and the friendship. Jones must choose between mutually exclusive options. Under some circumstances turning that kind of blind eye would be both immoral and illegal.

We are assuming again that whatever Jones does or does not do in deciding 'not to believe that Smith stole money from the firm' is actually possible; and that 'voluntary belief' and 'voluntary disbelief' refer to anything Jones can voluntarily do to increase the chances of having the belief he wants to have and of not having the belief he does not want to have, including behaving as if he holds the belief he wants to hold and does not hold the belief he does not want to hold.

Husband and wife

In Meiland's second story a wife finds what 'everyone would agree' is sufficient evidence of her husband's affair. But she believes the marriage is 'basically sound and can weather this storm'. Knowing she 'cannot conceal her suspicions' she 'decides to believe' her husband is faithful. Her husband 'eventually stops seeing the other woman, becomes more attentive to [her], and the marriage continues stronger than ever for many years'. Again Meiland does not think the wife should believe her husband was guilty just because there was sufficient evidence. As in the previous scenario believing against the evidence allows the believer to preserve 'some very precious things—a strong friendship, a good marriage' (1980:16).

With appropriate substitutions (3.7) becomes:

(3.7.1) {If anything is morally wrong, then} it is [morally] wrong of the wife to believe [knowingly or irresponsibly] that the husband is innocent of infidelity on insufficient evidence {in the absence of any conflicting and overriding moral imperative} [except if the unjustified believing is outside the wife's voluntary control].

The 'conflicting and overriding moral imperative' condition is perhaps even more

pertinent here, as the wife could have more at stake than Jones. But different circumstances could bring a different balance of imperatives. Assume it is conceptually and psychologically possible for the wife to ‘decide to believe’ her husband’s innocence, then factor in the redemptive power of forgiveness. In Meiland’s story she has blocked this path because there is nothing to forgive.

This is not a moral judgment. Her decision may be correct. But change some of the details and blocking the path to forgiveness might be the morally wrong thing to do, with the wrong of over-belief not trumped by another imperative. It is not that over-belief is only morally wrong if it blocks forgiveness, but that (following Clifford, minus his universality) over-belief is morally wrong in itself to at least some degree but, like all moral wrongs, can be trumped by other imperatives.⁴¹

Compare Meiland’s own verdict:

I believe that she has a duty to her husband, arising from their commitment to one another over a long period, to require a stronger basis for belief in his treachery than does, say, [a] private detective. ... It is not a matter of her having much evidence to the contrary. Instead, it is a matter of obligation toward someone to whom she has been very close. (1980:21)

It is understandable that because of the high stakes of their 15-year marriage it would be (morally) correct for her carefully to weigh up every consideration and conflicting imperative, leading to a decision to believe in his innocence. But in a different context it might lead to a decision to confront him, perhaps to establish the truth about their relationship. Her decision to believe in his innocence seems to amount to a decision to see the evidence in a particular light, to weigh it so as to support the belief she ‘decides’ to have.

Talk of ‘stakes’ anticipates a discussion in **Chapter 5** about whether such

⁴¹ ‘Think here of the spouse who, desperate to save a marriage, is willfully blind to all evidence of betrayal. Though belief in the straying spouse’s fidelity may be instrumental to keeping the marriage alive, self-deception of this sort seems hardly to be a cognitive virtue.’ (Taylor, 2007:151)

considerations can influence whether a belief is justified.⁴² We will therefore postpone further treatment until then, other than identifying another pertinent thread, that of the wife's (also 'high stakes') priority to do what she can to ensure their 'basically sound' marriage will 'weather this storm': 'her desire to continue the marriage turns out to be one of the factors that should influence what she does believe' (1980:22). In this respect her situation seems broadly similar to Jones's: intentionally doing whatever she can to increase the likelihood of having the belief she wants to have, including behaving as if she believes what she wants to believe.

So far neither of Meiland's cases seems problematic for EP4, principally because of its 'conflicting and overriding moral imperative' qualification.

This chapter's aim was to assess apparent objections to our principle, and if necessary modify it to accommodate them. We will resume discussion of religious belief and the criterion of sufficient evidence in **Chapter 4** and **Chapter 5** respectively, but apart from that we have covered what seem the major potential counter-examples, and finalised the wording of EP4. For completeness we will now briefly consider a few types of belief which seem to present less weighty challenges. The first two provide opportunities to clarify EP4's scope. The rest are potential objections that principles like EP4 and CP are too strict.

3.5 Personal preferences

Haack's critique of CP includes an 'apparently harmless unjustified belief' (1997, 2001:27): 'the apples I just selected are the best in the supermarket' (1997, 2001:24).

If I take a bite from a red apple followed by a bite from a green apple and believe the red is better than the green, would I need 'sufficient evidence' to justify my belief? Perhaps I believed the 'better' apple ranked, or would rank, higher when tested against independently testable criteria. If so that could qualify as an

⁴² p240ff. See also p295.

‘apparently harmless over-belief’, discussed later.⁴³ But we will assume ‘better’ and ‘best’ express purely personal preferences, regardless of any objective score or what anyone else thinks. It is not obvious how a belief like that can fall under EP4, or how it could possibly fail EP4, because in a purely personal preference it may be impossible to separate ‘evidence’ from ‘belief’. Preferring red to green is both ‘belief’ and ‘evidence’.

If a purely personal preference counts as a belief it would also arguably count as ‘prescriptive’: not morally prescriptive, but rationally or prudentially binding on my choices.⁴⁴ It would be irrational for me to choose green over red, unless in obedience to an imperative which I or someone else saw as overriding – to deny myself pleasure, or save the better apples for another person.

This kind of personal preference would not qualify as the kind of descriptive belief which is EP4’s intended scope. The same applies to the next category.

3.6 Aesthetic beliefs

As with personal preferences, and for similar reasons, EP4 should exclude evaluative aesthetic beliefs.⁴⁵ This does not mean excluding all beliefs about art. The belief that Degas was influenced by Japanese prints is a straightforward descriptive belief supportable by evidence.

But ‘Beethoven was the greatest composer of all time’ is not purely descriptive even though it might rest on one or more descriptive beliefs. Someone might believe Beethoven the greatest because of his orchestral ingenuity and his role in the transition from classical to romantic composition, both of which could be supported by evidence. But why this amounts to being the ‘greatest composer of all time’ is not something evidence could either support or falsify.

Arthur, who believes Beethoven was the greatest, would not be contradicting

⁴³ p121ff.

⁴⁴ This is different from a hypothetical prescriptive belief: see p44ff.

⁴⁵ As suggested earlier p14.

Beryl, who believes Wagner was. But Charles, who believes humans and chimpanzees last shared a common ancestor 4-6 million years ago, would be contradicting Diana, who does not believe humans and chimpanzees have a common ancestor.

Aesthetic beliefs represent a clear category easily distinguishable from the kind of descriptive beliefs EP4 applies to. The exclusion is not arbitrary and EP4 loses nothing by it. We will not however take the same line with this chapter's remaining beliefs.

3.7 Apparently harmless over-beliefs

We can tweak Haack's 'apparently harmless unjustified belief' to make it purely descriptive:

(3.8) The apples I just selected are the freshest in the supermarket.

'Freshest' here does not express personal taste but refers purely to the elapsed time from when they were picked. It is therefore the belief that no other apples in the supermarket were picked later than the ones the customer selected. In some circumstances 'there may be no great harm done by the mere belief' (1877, 1879:185). But in other circumstances there might be. A friend hearing me assert (3.8) and assuming I would not have believed this without adequate evidence⁴⁶ could be led to accuse the supermarket of misleading its customers, because he compared the date on my apples with the date on others I did not select. To count as a potential counter-example to EP4 the over-belief must be actually harmless, not just apparently harmless. Clifford does not think there are any such, however 'trifling and fragmentary'⁴⁷ they may seem:

Nor is that truly a belief at all which has not some influence upon the actions of him who holds it. (1877, 1879:181)

⁴⁶ See p23.

⁴⁷ See p15.

In **Chapter 1** we discussed Haack's objection that the 'risk of risk of harm' (1997, 2001:27) I run by making myself credulous is too remote for moral blame.⁴⁸ We countered this with Clifford's insistence that over-belief is insidious because like a 'plague' (1877, 1879:184) it can generate more of itself. We cannot know how harmless an individual over-belief is until all its direct and indirect consequences have materialised. This could stretch into the indefinite future, hence 'however seemingly trivial the belief, and however obscure the believer ... we have no choice but to extend our judgment to all cases of belief whatever' (1877, 1879:182).

So the category of apparently harmless over-beliefs does not seem to challenge EP4 because it is unlikely we could know in advance if an apparently harmless over-belief is really harmless. Adopting a principle like EP4 would therefore seem wise. If, in any concrete instance, we have reason to think it highly probable that a particular over-belief will prove harmless in the long run, then a 'conflicting and overriding moral imperative' could well tip the balance the other way. But this would not undermine EP4.⁴⁹

We will now consider potentially the most extreme variant of 'apparently harmless belief': where, assuming this is possible, the belief is completely private.

3.8 Private over-beliefs

For this category we must remove any risk to a third party. To succeed however I might have to make my every belief private, or follow different guidelines for private and what we will call 'active' beliefs.

An 'active' belief need not be asserted or voiced, and the believer 'may never have occasion to exhibit it in outward acts' (Clifford, 1877, 1879:185); but it (or something about how it was acquired) would or could influence the believer's behaviour (for example in the acquisition of other beliefs), which could then have

⁴⁸ p33.

⁴⁹ We are also assuming throughout that voluntary belief is possible, in the sense that 'voluntary belief' refers to all our voluntary acts and omissions aimed at increasing the chances of having the particular belief.

some direct or indirect impact on others, perhaps by influencing their own doxastic behaviour. Clifford thinks all our beliefs are active in this sense:

If a belief is not realized immediately in open deeds, it is stored up for the guidance of the future. It goes to make a part of that aggregate of beliefs which is the link between sensation and action at every moment of all our lives, and which is so organized and compacted together that no part of it can be isolated from the rest, but every new addition modifies the structure of the whole. (1877, 1879:181)

No belief like this is

a private matter which concerns [the believer] alone. Our lives are guided by that general conception of the course of things which has been created by society for social purposes. Our words, our phrases, our forms and processes and modes of thought, are common property, fashioned and perfected from age to age... (1877, 1879:181)

A private over-belief by contrast must have no possible impact outside the believer himself.

We will explore the options systematically to see if any meaningful category of private over-beliefs is problematic for EP4.

First consider Robinson Crusoe on his island, but with no Friday and no possibility of rescue. We can discount this scenario as EP4 presupposes a context of social and moral life.

In our second and third scenarios the agent lives within a community and has, respectively, either only private beliefs or a combination of private and active beliefs. The former would effectively exclude the agent from social interaction. He would be unable to communicate his beliefs, or communicate anything which derived directly or indirectly from his beliefs. Nor could anyone else know anything about any of his beliefs, about how he acquired them and what he based them on.

This would be another scenario ruled out by our assumed context of social and moral life.

In the third scenario the agent has both private and active beliefs, but fleshing this out presents difficulties. Beliefs where I could relax EP4 would not only be private, but could have no role in supporting any active beliefs. I would need separate networks of beliefs, because individual beliefs rarely if ever stand alone but are supported by other beliefs, which are in turn supported by further beliefs, and so on.⁵⁰ If an agent has a network of beliefs which he assumes are justified, and then finds one or more of his beliefs are unjustified, entire sections of the network could become suspect until the 'damage' is 'repaired' by removing and, if necessary, replacing the unjustified beliefs and/or restructuring the network.

For private over-belief to represent a meaningful counter-example to, or legitimate exclusion from, EP4, the agent would need to have (and be aware of having, and be able to maintain), separate belief networks to ensure an active belief never depended on a private belief, and no private belief ever contaminated an active belief.

But it is not just the private beliefs themselves which must not contaminate the active beliefs. Clifford's principal target is the agent's doxastic behaviour, the good or bad habits he cultivates as he acquires and sustains his beliefs. For private over-beliefs to evade EP4 the agent would need to apply consistently different criteria for acquiring and holding the two sets of beliefs. Even if this were possible, it is arguable whether the two sets of beliefs would qualify as the same kind of thing, which they would need to for private over-beliefs to represent a counter-example to EP4. Otherwise one could argue they would be outside EP4's intended domain of application because they are something significantly different.

We should also remember EP4 is intended as a moral principle for agents to follow: something agents could reasonably apply in practical circumstances. It is hard to see how it could be practical for an agent to keep his private beliefs, and

⁵⁰ See p81ff.

his criteria and methodologies for managing them, permanently insulated from all his other networked beliefs and doxastic behaviour. This requirement is imposed not by EP4 but by the attempt to articulate private over-beliefs as a counter-example to EP4.

The category of private beliefs was defined to evade Clifford's arguments, which presuppose a context of socially interactive 'normality' where people have, share and assert beliefs among other interactions with different types and degrees of moral significance. Our conclusion is that, even assuming a possible mode of 'believing' disconnected from the rest of social interaction, the category of private over-beliefs could claim at best a hollow victory. They would evade culpability only by being trapped inside a bubble of existence devoid of moral value.

Chapter 4 will revisit private over-beliefs in the context of religious belief.⁵¹ The present chapter's final category is of cases where although we have neglected our doxastic duties the resultant beliefs are actually true.

3.9 Unjustified true beliefs

We have already encountered this kind of belief under **Testimony**,⁵² and in a variant of Clifford's opening parable.⁵³

[S]uppose that the ship was not unsound after all; that she made her voyage safely, and many others after it. Will that diminish the guilt of her owner? Not one jot. When an action is once done, it is right or wrong for ever; no accidental failure of its good or evil fruits can possibly alter that. The man would not have been innocent, he would only have been not found out. The question of right or wrong has to do with the origin of his belief, not the matter of it; ... not whether it turned out to be true or false, but whether he had a right to believe on such evidence as was before him. (1877, 1879:178)

⁵¹ Eg p154.

⁵² p91.

⁵³ p14.

Clifford's ship-owner let himself 'believe for unworthy reasons', without 'doubting' or 'judicially and fairly weighing evidence'. Although ostensibly what makes an over-belief undesirable is the consequences of its falsity if it is false, all a believer can control is cultivating good doxastic habits and avoiding bad ones – for their direct and indirect effects on himself and others. If he is sometimes lucky because, through no contribution or care on his part, his belief 'may be true after all' (1877, 1879:185), that is irrelevant. He could apply the same policy in the future and sooner or later believe something false with unfortunate consequences. And he would be guilty, in however small a way, of 'doing this great wrong towards Man, that I make myself credulous'.

There is no obvious reason to regard these over-beliefs as problematic for EP4, or exempt them from it. We will revisit them in **Chapter 4**,⁵⁴ as James considers them relevant for religious belief.

3.10 Conclusions

This chapter has discussed types of belief which, when an evidence principle is applied to them, could appear to give challenging results. Some, like involuntary and foundational beliefs, could threaten to derail the principle. Or if beliefs about the past or future, beliefs based on testimony, and 'positive' believing linked to personal or interpersonal aspiration are all ruled out, legitimate beliefs could shrink to a paltry residue.

But far from threatening our principle, these challenges have helped clarify its scope and spirit, evolving it into its final version as EP4. In the process five key assumptions (3.9.1–5 below) have emerged or reappeared, helping to flesh out the intention behind EP4. We have used different combinations of these assumptions when dealing with most of the challenges. The exceptions are personal preferences and aesthetic beliefs, which we placed outside EP4's intended scope.

⁵⁴ Eg p143 footnote.

- (3.9.1) EP4 is a moral principle. This is implicit in CP and explicit in how Clifford presents CP. EP4 is furthermore purely a self-standing moral principle applying to agents' doxastic behaviour. It is not separately intended as a comprehensive criterion of epistemic justification derived from moral considerations. Assumptions (3.9.2–4) either follow from this or unpack more of how EP4 behaves as a moral principle.
- (3.9.2) EP4's principal target is our voluntary acts and omissions relating to how we acquire, evaluate, hold and reject our beliefs, even if these only indirectly influence what beliefs we end up with. This also parallels how Clifford presents CP.
- (3.9.3) EP4 presupposes a moral context which is inherently social. This is also implicit in CP as Clifford's principal rationale is the social and moral consequences of what he sees as good and bad doxastic habits. The strength of this assumption for EP4 is that beliefs and doxastic behaviours which social and moral life depend on cannot be illegitimate.
- (3.9.4) A moral principle like EP4 is not absolute in relation to other moral principles: at times it will be right to transgress EP4 if another imperative overrides it. We see this as implicit in what a moral principle is, although Clifford does not seem to share it.
- (3.9.5) Beliefs are not atomic but embedded in interconnected networks of other beliefs. This is clearly related to (3.9.1): EP4 primarily targets the agent's doxastic behaviour, and it seems a fact about human believing that an agent's beliefs are interconnected in networks rather than atomic. However (3.9.5) is identified as an independent assumption.

In some of his anti-sceptical defences, for example when justifying 'the assumption that nature is practically uniform as far as we are concerned' Clifford employs the conception of belief as a guide to action: belief is 'no guide to action until it takes to itself' what 'goes beyond experience', which means 'assuming that

what we do not know is like what we do know' (1877, 1879:206). In respect of EP4 we need not commit to an independent assumption that belief is a guide to action, although it would seem a plausible consequence of (3.9.3). Social and moral life would be impossible if beliefs could not serve as guides to action.

In no belief category reviewed in this chapter have we yet found a reason to think EP4 unsound. It permits the beliefs we would expect it to permit and forbids the beliefs we would expect it to forbid. But we have not yet considered religious belief, a key battleground between Clifford and James. Even if Clifford does not use CP to undermine religious belief explicitly,⁵⁵ James presents 'The Will to Believe' as a counterattacking 'defence of our right to adopt a believing attitude in religious matters' (1896, 2000:198). Reason enough to devote the entire next chapter.

⁵⁵ It is worth noting however that Clifford originally presented 'The Ethics of Belief' (1877, 1879) in April 1876 as an address to the Metaphysical Society, of which he was a member. The context was an extended debate involving James Fitzjames Stephen, Thomas Huxley, Shadworth Hodgson and others about belief in miracles and other supernatural events (Madigan, 2009:73). The Metaphysical Society was originally set up out of a concern shared by many of its members (but not Clifford himself) that declining religious faith among the educated elite could lead to general deterioration in social morality (Madigan, 2009:1).

Chapter 4 | William James and the believing attitude

4.1 Introduction

We have considered several challenges to an evidence principle covering descriptive beliefs. Most we rejected. A few caused us to adjust it towards its eventual form as EP4. Religious belief is different though. While James sees it as reason to reject an evidence principle, a freethinker might see it as what an evidence principle is there for.¹

This chapter has therefore two objectives. One is to see if anything in James's defence of religious belief threatens EP4. There is no intention to reconcile James's position with EP4, more to see if any of his arguments might persuade us that it would be more reasonable to allow a belief which EP4 might disallow. The other objective is to identify potential advantages, specifically in relation to religious belief, of adopting a principle like EP4. Pursuing both we will unpack ramifications of James's 'justification of faith' (1896, 2000:198), both in themselves and in contrast with EP4. We aim to show that EP4 should survive any Jamesian counter-attack, and that if there is a sound Jamesian position on religious belief which is incompatible with EP4, then the choice between them is ultimately a moral one.

We first evaluate the combined descriptive and prescriptive content of a typical religious belief against EP4² (introducing the first potential benefit of EP4³) and then review James's conceptual framework,⁴ including his treatment of Pascal's Wager.⁵ Next we analyse James's own 'religious hypothesis' (1896, 2000:214)⁶ into descriptive and prescriptive components. Taken together these construct a 'genuine option'⁷ which James sees each individual having the right to decide for

¹ See for example Singer (2005).

² pp131–140.

³ p137ff.

⁴ pp140–165.

⁵ p145ff.

⁶ p165ff.

⁷ p167ff.

herself according to her own risk appetite for 'knowing the truth' versus 'avoiding error'.⁸

We then elicit further potential negative consequences of the kind of religious over-belief James would allow,⁹ and consider if the risk of these can at least sometimes (as James would insist) or never (as Clifford would maintain) impact just the believer herself. EP4 does not outlaw religious over-belief completely. It permits it if it is the only route to a net moral good; but it would then be obligatory. James though allows someone either to believe or disbelieve an undecidable hypothesis within a 'genuine option'. We ask if this would permit someone to cede her moral sovereignty to God,¹⁰ and whether this could contradict James's own ethic of tolerance.

We then offer a rather more generous interpretation of James's overall position¹¹ and suggest two 'Jamesian' moral principles as potential alternatives to EP4.¹² A Jamesian and a proponent of EP4 might well make similar judgments in practice;¹³ but a difference of principle still remains, reflecting Jamesian liberalism versus Cliffordian social protectionism, although these may prove difficult to disentangle in practice.

We conclude that EP4 has little to fear from James's 'justification of faith'.¹⁴ At most it offers a contrasting moral individualism which seems hard to justify as a reason for rejecting EP4.

⁸ p163ff.

⁹ pp177–190.

¹⁰ p185ff.

¹¹ pp190–195.

¹² p195ff.

¹³ p200.

¹⁴ p203ff.

4.2 Compound structure of religious belief

EP4 and prescriptive beliefs

Chapter 2 concluded that applying an evidence principle to prescriptive beliefs seems problematic. That might suggest EP4 would permit any prescriptive belief. But other constraints apply. Combinations of prescriptive and descriptive beliefs should be consistent, so it would be incoherent to believe you must spin straw into gold if your descriptive beliefs say this is impossible.

Our prescriptive beliefs should also be mutually consistent. It would be wrong to believe 'you should not lie' and 'you need not tell the truth' simultaneously. So although EP4 only covers descriptive beliefs that does not make it powerless against every prescriptive belief. As EP4 is itself prescriptive it cannot be held alongside another incompatible prescriptive belief. Holding EP4 would for example rule out thinking one should believe God exists even without sufficient evidence or a conflicting and overriding moral imperative. As this is a plausible reading of the First Commandment, EP4 would conflict with at least some imperatives associated with religious belief, and we would be faced with a moral choice as to which to adopt. With all this in mind, we will now consider the structure of religious belief.

Descriptive and prescriptive components

This chapter will eventually articulate three potential benefits, specifically in relation to religious belief, of adopting a principle like EP4.¹⁵ All three relate to how components of religious belief can link together. Below is a familiar kind of compound religious belief which James would seem to allow, CP would probably disallow, and EP4 would only permit under strict conditions. It has both descriptive and prescriptive content:

(4.1.1) God exists. [Descriptive.]

¹⁵ See p137, p178 and p181 respectively.

(4.1.2) God is the source of morality. [Descriptive, strictly speaking, if only as a facet of ‘God is the creator of all things’.¹⁶ It could however be considered to have a prescriptive implication, separated out as (4.1.3) below.]

(4.1.3) What God commands must be obeyed. [Prescriptive.]

This might seem the obvious point that for many people religious belief is associated with specific moral views. But its very obviousness can disguise the unique ways these components can link together, which are significant for a principle like EP4.

On the face of it the purely descriptive content – ostensibly the only part relevant to EP4 – appears harmless. Why should someone’s belief in God’s existence (4.1.1) matter? Even James saw the ‘debate between materialism and theism’ as ‘idle and insignificant’ (1907, 2000:47) – when we look backwards:

It makes not a single jot of difference so far as the past of the world goes, whether we deem it to have been the work of matter or ... a divine spirit[.]
(1907, 2000:45)

It is only when we ask ‘what does the world *promise?*’ (1907, 2000:48) that theism and materialism point to ‘wholly different outlooks of experience’ (1907, 2000:49). For James belief in God satisfies the need for an ‘eternal moral order’ (1907, 2000:50), which we can choose whether or not to align ourselves to by ‘getting upon the winning side’ (1896, 2000:216). What matters is the impact on behaviour; and the prescriptive content of religious belief is what directs people’s actions and shapes how they treat each other.

Believing (4.1.1) does not necessarily mean believing (4.1.2) and/or (4.1.3).¹⁷ Kant for example thought the route to the moral law was conceptually independent of religious belief but that faith in God represents a ‘need of reason in its practical

¹⁶ See p135.

¹⁷ See however ‘Reading (ii)’ below.

use' to 'presuppose the existence of God not only if we *want* to judge, but because we *have to judge*' (1786:12). But Kant notwithstanding, (4.1.2–3) together describe a frequent and familiar divine attribute. If you believe a God exists (4.1.1) whose imperatives (4.1.2) must be obeyed (4.1.3) then you believe something about what is and something about what ought to be. The parts appear logically distinct, but someone holding compound belief (4.1.1–3) cannot stop believing one part without affecting the other parts.

Compare:

(4.1.1BJ) Boris Johnson is the UK prime minister.

(4.1.2BJ) We should take seriously everything Boris Johnson says.

(4.1.1BJ) and (4.1.2BJ) can be believed or disbelieved independently of each other. But believing (4.1.1) without (4.1.2–3) would be believing in a different kind of God than someone who believed all of (4.1.1–3) would believe in. And someone believing (4.1.3) without (4.1.1) – assuming this were possible – would not be bound the way someone believing all of (4.1.1–3) would be.

There are also multiple ways of believing (4.1.1–3), which to an extent reflect different perceived distributions of contingent and necessary relationships between their various components. Consider for example two readings of Exodus. Reading (i)¹⁸ sees God as a character in an account which could be either historical or fictional. In a sequence of purely contingent story elements God speaks to Moses, inflicts plagues on Egypt, then gives Moses tablets of stone on Mount Sinai. In theory God could have chosen not to deliver the Commandments, or picked a different time and place. Reading (ii) by contrast sees the time-structured Exodus narrative as an imperfect attempt at articulating God's timeless features: the Decalogue's prescriptivity would be an essential aspect of what

¹⁸ We will be considering a number of sets of options, each with its own labeling: readings (i) and (ii); believers A and B; and views α , β , γ and δ . This is to indicate that they are separate sets, although some will overlap: for example believer B may have view δ .

someone who believes in God believes is true of God, including God's role in how morality enters the world.

Now consider extremes of belief in relation to each reading. Believer A subscribes to (i) and not (ii). She also believes the Exodus events actually happened, that God existed then and exists now (4.1.1). Her belief involves specific moral imperatives, and the way those imperatives connect with her belief in God (4.1.1) can be broken into steps. She believes her God is the God who spoke to Moses, and she has decided (implicitly or explicitly) to obey this God. She sees herself bound by the Ten Commandments (4.1.3) because her God gave them to Moses (4.1.2).

Believer B subscribes to reading (ii) and not (i). She sees the Exodus story as an allegory which is literally false. Key to her conception of God is that the descriptive aspects – (4.1.1), (4.1.2?) – and prescriptive aspects – (4.1.2?), (4.1.3) – are inseparable. She believes God's role as the source of morality and the binding nature of God's commandments are necessary attributes of God. She believes God exists, either necessarily or contingently.¹⁹

We will now consider different views of how (4.1.2) and (4.1.3) relate together.

View α might see (4.1.2–3) as equivalent to:

(4.1.2) God is the source of morality... [Premise]

(4.1.5) ...**therefore** God's commands must be obeyed. [Conclusion]

View β though might think a step is missing between (4.1.2) and (4.1.5):

(4.1.4) If X is the source of morality, then whatever X commands must be obeyed. [Additional premise]

¹⁹ Whether B believes God's existence is necessary or contingent is not relevant in this context.

View α would see (4.1.4) as self-evident and redundant. Perhaps α thinks that if God is omniscient and omnipotent and the creator of all things, therefore God created morality, because 'all things' must include morality. Therefore God is the source of morality. Therefore God's commands must be obeyed. This gets from an 'is' to an 'ought' apparently seamlessly.

The step β thinks is missing (and needing independent assent) is where you place yourself in God's thrall (4.1.4). Just because God is omnipotent and the creator of all things does not mean you must be in God's power. Or β might reject the idea of an external 'creator of morality', because if morality is something you must obey, it cannot be something an entity independent of you can create.

View γ might however be that (4.1.2–3) together indicate an 'institutional' context similar to that of 'promising' in Searle's (1964:44) attempt to derive an 'ought' from an 'is'. This could be summarised as:

(4.2.1) *Jones promised to pay Smith 5 dollars.* [Searle's original (2)]

(4.2.3) [Therefore] *Jones ought to pay Smith 5 dollars.* [Searle's original (5)]

Searle adds a number of intervening premises, but the key one, equivalent to (4.1.4), is 'tautological' and added 'for the purpose of formal neatness' (1964:46):

(4.2.2) *All promises are acts of placing oneself under (undertaking) an obligation to do the thing promised.* [Searle's original (2a)]

The institution of promising (4.2.2) legitimises the move from 'is' (4.2.1) to 'ought' (4.2.3). Jones and Smith are interacting within a social context where promising is an available transaction. On that assumption (among others), because Jones promised to pay Smith \$5 (4.2.1) we can derive the prescription (4.2.3) that Jones ought to pay Smith \$5.

View γ might therefore see God as an entity playing a quasi-'institutional' role in relation to other entities (particularly conscious ones with free will). But, assuming

(4.1.2–3) alone do not imply the God's existence, then (4.1.2–3) confer no obligation either, as without (4.1.1) the God, with its role and its powers, remains just a possibility. Like (4.2.1), (4.1.1) calls the obligation into existence. Combined with the assumed is-to-ought transition of (4.1.2–3), (4.1.1) generates binding prescriptive content from the descriptive assertion that a specifically defined entity exists. Sincerely asserting (or believing) the God exists (4.1.1), where 'God' incorporates the content of (4.1.2–3), is a commitment which parallels the commitment Jones makes in (4.2.1), and with similar implications for obligation.

But just as lying to the Gestapo can be simultaneously an instance of lying, of courage, and of loyalty,²⁰ believing a God exists (4.1.1) – with 'God' characterised by (4.1.2–3) – can be both a belief about what is the case and a commitment implying obligation.

Yet another view (δ) about how (4.1.2) and (4.1.3) relate could be that of believer B, for whom the kind of analysis just offered is anathema as it falsifies the essence of the God she believes exists (4.1.1). View δ might see both (4.1.2) and (4.1.3) contained within, and therefore entailed by, (4.1.1). How true or even coherent such a view might be is less important than whether there are people who hold it (apparently yes); and whether it qualifies for the kind of defence James will provide.

There are doubtless other views – perhaps as many as there are believers holding versions of (4.1.1–3). Whatever their individual differences, the structured combination of (4.1.1–3) can have significant practical and moral consequences. Adding content about what the God forbids and demands – that, say, believers should cede their moral sovereignty to the God and persecute blasphemers – can lead to repression. On its own the belief that the God demands persecution of blasphemers would be descriptive, but combined with (4.1.1–3) generates the prescriptive belief that blasphemers should be persecuted.

²⁰ See p113.

This brings us to the first potential benefit of a principle like EP4 in relation to religious belief: that it usefully delegitimises the kind of religious belief which includes repressive content and is logically structured to make it binding, regardless of whether the believer herself is aware of this. It need only target the descriptive component, because that is where the problem lies. The issue is not the different prescriptive beliefs which different people and communities hold, but the descriptive beliefs about the existence of entities which ordain those imperatives and make them binding.

This might seem counter-intuitive, as the commandment to persecute blasphemers is what is repressive. Without that being held and acted on blasphemers will escape persecution. But the danger is not just of holding an imperative which others might condemn. Moral persuasion will be appropriate and necessary as long as societies survive. The specific danger is of holding the imperative as an ineluctable commandment from a being believed to be uniquely responsible for creating and sustaining all there is, including those so blind they do not share the belief. Moral persuasion can struggle against beliefs about such 'great cosmical matters' (James, 1896, 2000:214).

This would be a reason to advocate an evidence principle applied to descriptive beliefs,²¹ demanding evidence that the God-as-source-of-moral-law exists. In the absence of evidence, or of any conflicting and overriding moral imperatives, it would be morally wrong to believe the God exists. The God could not itself be the source of those other imperatives, at least in respect of the moral choice concerning the descriptive belief in its existence. Even if the God existed and was the source of all moral imperatives, EP4 would withhold the right to believe, without sufficient evidence, either in its existence or that it was the source of all moral imperatives.²²

²¹ See p41.

²² If God did exist, and was the source of all moral imperatives, then assuming EP4 was itself a sound moral imperative, God would be its source. It could be considered illogical for a God to have ordained EP4. If so, then the existence of God as source of moral imperatives might rule out EP4. But we do not need to know the source of a moral imperative in order to apply it or consider ourselves bound by it. So if we have no reason to reject EP4 and no evidence for God's existence, we would be free to hold EP4.

The promotion of EP4 here in respect of religious belief is therefore supported by empirical claims. These include the more general ones Clifford uses to support CP, articulated earlier as (1.1.1–13).²³ In the case of EP4 and religious belief we would add further claims relating to our understanding of pertinent is–ought transitions, and our observation of them in practice. Above we sketched a number of different views (α – δ) as to how the transition from descriptive belief (4.1.1–2) to prescriptive belief (4.1.3) proceeds. Which if any of these is ‘correct’ is less important than the fact that a variety of such views are actually held; and none seems obviously incoherent at first sight, especially given the particular content of descriptive religious belief, rich in superlatives, universals and supernatural elements:

A religious community can be defined in terms of the grip of imaginings that sprinkle fairy dust on the transition between “is” and “ought,” ensuring communal support for whatever transition has become salient, enforcing uniformity, and making dissent difficult or impossible. (Blackburn, 2007:189)

A critic might now remind us of ‘sophisticated’ religious believers for whom, say, God cannot ordain immoral commands, or for whom God provides the motivation and/or obligation to adhere to moral imperatives, but not necessarily those moral imperatives themselves – or at least not single-handedly. While an unsophisticated believer may cling to commandments from ancient texts and/or disseminated by priests claiming a direct line to their shared God, a sophisticated believer might see the imperative of working out what God commands as both far more difficult and far more important.

But the sophisticated believer is still answerable to Clifford’s charge about the ‘greater and wider evil’ which ‘arises when the credulous character is maintained and supported’, the risk to society being that it should ‘become credulous, and lose the habit of testing things and inquiring into them; for then it must sink back into savagery’ (1877, 1879:185-6).²⁴ In a social context sophisticated over-belief risks encouraging unsophisticated over-belief, and the content of sophisticated

²³ See pp18–21.

²⁴ See pp14–17.

over-belief cannot exculpate it – unless that content is such that the intended result of the over-belief provides the kind of conflicting and overriding moral imperative which EP4 accommodates.²⁵

In this section we have made a case for EP4 specifically in relation to religious belief by arguing as follows:

- (i) Religious belief typically contains both descriptive and prescriptive content.
- (ii) The two types of content are linked in ways particularly significant for an evidence principle like EP4:
 - (a) The descriptive beliefs make the prescriptive beliefs binding. Delegitimising the descriptive beliefs would undermine the prescriptive beliefs.

This assumes the believer holds the entire relevant set of descriptive and prescriptive beliefs. A further feature of the coupling between the two types of component can make this more likely:

- (b) The cosmic and/or cosmological content of some of the descriptive claims can obscure the transition between descriptive and prescriptive components,²⁶ articulated here as (4.1.1–3). The transition can then lend itself to a variety of interpretations, including direct entailment.
- (iii) In view of the combined descriptive and prescriptive components and the coupling between them, a principle like EP4 targeting the

²⁵ We will eventually (p185ff) make a similar point in relation to James, whose own 'religious hypothesis' can be analysed into descriptive and prescriptive content (p166ff).

²⁶ See also p176ff.

descriptive content would be beneficial against potential negative consequences of the prescriptive content.

Conclusion (iii) does not outlaw religious belief entirely. In a specific context believing in a God might, not necessarily for Kantian reasons,²⁷ bring otherwise unachievable net moral benefits. It would then be morally right (therefore permissible) to believe in the God because this would trump the wrong of over-belief. But the onus would be on the believer to establish independently that this was so. The God's commandments could not justify her decision.

With all this in place we will now evaluate James's counterattack against CP in relation to EP4. We will see if he can legitimise religious over-belief even in the absence of independent moral justification, and whether his own 'religious hypothesis' avoids the negative consequences we have begun to identify. We start by discussing his belief framework, including his treatment of Pascal's Wager.

4.3 James's belief framework

Contrast with Clifford

The central thesis of James's 'The Will to Believe' is that

W2B *Our passionate nature not only lawfully may, but must, decide an option between propositions, whenever it is a **genuine option that cannot by its nature be decided on intellectual grounds**; for to say, under such circumstances, "Do not decide, but leave the question open," is itself a passionate decision,—just like deciding yes or no,—and is attended with the same risk of losing the truth. [Emphasis added] (1896, 2000:205)*

Although W2B covers any 'genuine option', religious belief is James's main focus. In W2B we see an ethics of belief in clear conflict with CP and EP4, as there are

²⁷ p132.

kinds of belief which both would outlaw but James would permit. As previously mentioned however²⁸ we aim to show that if there is a sound Jamesian position on religious belief which is incompatible with EP4, any choice between them would ultimately be a moral one.

Clifford and EP4 would deem an individual morally wrong, other things being equal, to hold a descriptive religious belief without evidence. But James would allow her either to believe or disbelieve, although she must do one or the other. Clifford and James would agree a third option exists, in theory: to leave the question open. But for James leaving the question open and deciding not to believe are for practical purposes identical. So where there is no evidence either way, if it is legitimate to incur the practical consequences of non-belief (by disbelieving or by avoiding the decision), then incurring the practical consequences of belief is also legitimate.

We should note the following contrast. For Clifford:

(4.3.1) X is permitted to believe that p if she has sufficient evidence that p .

Clifford is not saying X must believe that p if she has sufficient evidence that p . But if she has insufficient evidence that $\sim p$ (which would be so if she has sufficient evidence for p) then she should not believe that $\sim p$. So if she has sufficient evidence that p , although she is permitted not to believe either p or $\sim p$, she is only permitted to believe that p .

James would accept (4.3.1) but add, from W2B:

(4.3.2) X is permitted to believe that p or that $\sim p$ if her belief proceeds from a 'genuine option' which cannot be decided on intellectual grounds.

We can reasonably assume James's 'intellectual grounds' include or are identical with 'evidential grounds', as he is arguing against CP which talks of 'insufficient

²⁸ See p129.

evidence'. So to understand what is at issue between James and Clifford we must unpack 'genuine option'.

Genuine option

James's 'genuine option' rests on 'some technical distinctions' (1896, 2000:198). These form part of his purely descriptive analysis of belief, which underpins his more normative claims.²⁹

According to James a 'hypothesis' is 'anything that may be proposed to our belief' (1896, 2000:199), and may be either 'live' or 'dead'. A live hypothesis 'appeals as a real possibility to him to whom it is proposed' and makes an 'electric connection with your nature', while a dead hypothesis 'refuses to scintillate with any credibility at all'.³⁰ The same hypothesis may be live for one person but dead for another.

An 'option' is a 'decision between two hypotheses', and may be:

- (i) **Living** or **dead**
- (ii) **Forced** or **avoidable**
- (iii) **Momentous** or **trivial**

A **living** option (i) is a choice between two live hypotheses. James suggests theosophy versus Islam as a **dead** option for his lecture audience, whereas agnosticism versus Christianity might be living. (He does not explicitly allow an option between more than two hypotheses, but it seems reasonable to assume his schema permits it.)

A **forced** option' (ii) is a 'dilemma based on a complete logical disjunction, with no possibility of not choosing': for example 'Either accept this truth or go without it' (1896, 2000:199). An option is not forced if it can be avoided by not choosing.

²⁹ The 'normative' discussion begins on p160.

³⁰ See however **Sequence of believing**, p144ff.

'Either call my theory true or call it false' is **avoidable** by not judging either way. James speaks as if the forced/avoidable distinction, unlike the live/dead distinction, is intrinsic to the option, not relative to the individual. '[N]o standing place outside of the alternative' (1896, 2000:199) suggests it applies to everyone. But we will question the practical significance of this later.³¹

A unique, once-in-a-lifetime opportunity would be **momentous** (iii). His paradigm case is whether or not to believe the 'religious hypothesis' (1896, 2000:215), which could represent one's 'sole chance in life of getting upon the winning side' (1896, 2000:216).³² **Trivial** options however 'abound in the scientific life' and are 'when the opportunity is not unique, when the stake is insignificant, or when the decision is reversible if it later prove unwise' (1896, 2000:200). We could reasonably assume momentousness versus triviality is also relative to the individual: one might see a stake as significant which another thinks insignificant.³³

James's examples of momentous options are mostly actions or opportunities rather than beliefs. His key secular example is choosing to decline or accept an invitation to join Nansen's polar expedition, which

would either exclude you from the North Pole sort of immortality altogether or put at least the chance of it into your hands. (1896, 2000:200)

Someone turning this 'unique opportunity' down 'loses the prize as surely as if he tried and failed'. This is hard to construe in terms of hypotheses 'proposed to our belief'. The slide between 'belief' and 'action' is not unique to the momentous/trivial dichotomy. Avoidable options include 'Choose between going out with your umbrella or without it' and 'Either love me or hate me' (1896, 2000:199). Others have both belief and behavioural dimensions: being Muslim or Christian is partly about what you believe and partly about what you do.

³¹ p168ff.

³² When and if the believer gets to that 'winning side' her belief would presumably be revealed as having previously been an 'unjustified true belief': see p125ff.

³³ In theory uniqueness and irreversibility might seem objective rather than subjective criteria but in practice they would contribute little to the 'momentousness' of an option in the absence of significance. They might also be subject to agent-relative criteria as to what counts as unique or as irreversible.

The indistinct boundary may be a by-product of James's pragmatism,³⁴ or merely indicate that genuine options can involve both potential beliefs and potential actions. In either context a **genuine** option is **forced, living and momentous**.

Sequence of believing

If a 'hypothesis' is 'anything that may be proposed to our belief' then it precedes the belief, as a candidate belief. It can be a hypothesis before being a belief, because its liveness for an individual is a precondition for her believing it. James is not claiming the liveness justifies the belief, but that it plays a part in giving rise to the belief.³⁵ He seems to claim it as an independent contingent psychological fact that a hypothesis can and must be live for someone before she can believe it,³⁶ with the liveness accounting for why she even considers believing it. (This would qualify as an empirical claim, like those we identified behind Clifford's case for CP.³⁷)

James introduces 'hypothesis' in terms of belief, not vice versa. If a belief starts as a hypothesis, the hypothesis must have been live. Could someone see or hear something and immediately form a belief? There seems no need to interpolate a 'hypothesis' into this sequence unless it was obviously there, and we have no reason to think James would disagree. So the notion of a 'live hypothesis' may be needed to explain some but not all beliefs. His exception (4.3.2) would then only apply to hypotheses we in some way choose to believe (assuming this is possible), but not to involuntary beliefs (assuming they are possible).

The paradigm case of over-belief James exculpates is therefore a *voluntary* choice between live hypotheses constituting an undecidable genuine option. He says nothing about permitting an over-belief acquired outside a genuine option, so perhaps more passively and involuntarily. EP4 on the other hand, following

³⁴ See **Pascal's Wager** p147 and **4.4 James's case for religious belief** p167ff.

³⁵ There could be a parallel between James's 'live hypothesis' and the Bayesian notion of non-zero personal probability: see p213ff.

³⁶ But see next paragraph.

³⁷ p18.

Clifford, applies the voluntary/involuntary distinction primarily in relation to doxastic *habits*:³⁸ the more the over-belief results (even indirectly) from (voluntary) doxastic negligence the more culpable the believer is; whereas the more the over-belief is outside the believer's voluntary control, and therefore *involuntarily* acquired and/or held, the less culpable she is. We are not claiming James would either agree or disagree, just that his focus is a category of voluntary over-beliefs he thinks we are entitled to hold.

Before defending his own position James discusses Pascal's Wager (Pascal, 1966:149ff) to consider how belief can involve both 'will' – 'our passionate and volitional nature' (1896, 2000:200) – and 'intellect'. Although James's defence of religious belief also involves a 'live hypothesis', he first uses it to augment Pascal's rather different stance. This will be an opportunity to deepen our understanding of James's 'live hypothesis', and test EP4 against his interpretation of the Wager.

Pascal's Wager

While James claims religious over-belief is *permissible*, ostensibly on *moral* grounds, Pascal sees it as *obligatory*, on *rational, pragmatic* grounds. EP4 by comparison judges religious over-belief *impermissible*, other things being equal, on *moral* grounds. So Pascal's Wager, either with or without James's help, could challenge EP4.

Pascal applies decision theory to argue that believing a God exists is in an individual's best interests. Assume the individual is convinced that if she decides to believe she could win an infinite benefit (if the God exists) or incur a finite loss (if not); whereas if she decides not to believe she could win a finite benefit (if the God does not exist) or incur an infinite loss (if it does).³⁹ As long as she assumes a non-zero probability that the God exists, then the 'expected value' (value of outcome times probability) of the combination 'believe in the God's existence' plus

³⁸ See p74ff.

³⁹ We are assuming that if God does exist but the punter wagers not to believe, she will then, as one of the 'damned', suffer a 'wretchedness' to parallel the 'infinity of infinitely happy life to be won' if she had wagered to believe (Pascal, 1966:149-151).

'the God exists' outweighs the expected value of any other possible combination. Her rational choice is therefore to believe the God exists.

One potential problem is the leap from being convinced that acting in a certain way is in your best interests (where 'acting' includes believing) to being convinced that something is true. Surely pragmatic reasons cannot bear on the truth of a hypothesis? Pascal acknowledges this doubt, advising someone convinced that believing in God is the right thing to do but 'so made that [she] cannot believe', to

learn from those who ... have been cured of the affliction of which you wish to be cured... They behaved just as if they did believe, taking holy water, having masses said, and so on. That will make you believe quite naturally... (1966:152)⁴⁰

Hacking agrees:

Pascal ... accepts as a piece of human nature that belief is catching⁴¹]... The two possible acts are not, "Believe in God" and "Do not believe." One cannot decide to believe in God. One can decide to act so that one will very probably come to believe in God. Pascal calls that the wager that God is. To wager that He is not is to stop bothering about such things. (1972:188)

Franklin adds that Pascal

does not even believe one could come to believe in God through having perfect evidence, let alone through a desire to believe. Faith ... is a free and undeserved gift of God[:] ...what the Wager motivates is not belief, but action, 'saying masses and the rest', which will then dispose the seeker to receive God's grace. (1998:110)

⁴⁰ This reads like Alston's category III (if not category II) doxastic voluntarism: see p77ff.

⁴¹ Clifford would agree: 'pestilence'; 'plague' (1877, 1879:184).

On this understanding, Pascal's argument provides a reason to hold a prescriptive rather than descriptive belief.⁴² Persuading someone to behave in a particular way can involve persuading her she ought to behave like that. This 'ought' need not be a moral ought. The Wager's payoff matrix appeals to an individual's interest in maximising personal utility, so if successful would result in a purely prudential 'ought'. We will revisit this later.⁴³

If the outcome of a pragmatic argument is a prescriptive belief, then acquiring it falls outside EP4. But the prescriptive belief here promotes behaviour likely to facilitate acquiring a descriptive over-belief. While EP4 might not stop you adopting, on pragmatic rather than evidential grounds, a prescriptive belief (that you should indulge in masses and holy water), it would stop you acting on it if, as Pascal maintains, it could facilitate acquiring a descriptive over-belief. Behaviour Pascal advocates as prudentially obligatory EP4 would consider morally wrong.⁴⁴

So even if Pascal aimed to persuade us to undertake a particular course of action, it is clear that the intended result of his Wager is a descriptive belief in a God's existence. We will therefore continue to see the Wager as providing an argument for a descriptive belief, remembering this is an abbreviated way of talking about all the steps meant to induce belief that Pascal describes.

James on Pascal

James describes the Wager as an attempt to 'force us into Christianity by reasoning as if our concern with truth resembled our concern with the stakes in a game of chance'⁴⁵ (1896, 2000:200-201). He thinks a decision to believe in God 'adopted wilfully after such a mechanical calculation' lacks the 'inner soul of faith's reality' (1896, 2000:201):

⁴² See p14.

⁴³ See p153.

⁴⁴ Whether someone finds it psychologically possible to believe something, directly or indirectly, on the basis of a pragmatic argument is distinct from whether she would be justified in doing so. Our concern is with the latter.

⁴⁵ He could be describing Bayesianism: see p213ff.

unless there be some pre-existing tendency to believe in masses and holy water, the option offered to the will by Pascal is not a living option.

This suggests that whatever makes something a living or dead option for us, it is not volition ('will'). The existence of God (specifically the Catholic God) must have been a live hypothesis for the individual to be persuaded to try masses and holy water.

There is not space to do full justice to Pascal's arguments, but we should assess James's claim that to be convincing the Wager needs buttressing with a prior live hypothesis to defend it against the 'many Gods' objection.⁴⁶

The Wager assumes the God, if it exists, provides most of the benefits and costs: rewarding believers with infinite bliss and non-believers with infinite wretchedness. As long as the individual has only the live hypothesis that a specific God exists plus the live hypothesis that that God does not exist, the Wager can claim to provide a compelling pragmatic reason for choosing to act so as to maximise the probability of believing in that God rather than choosing not to. But if she has simultaneous live hypotheses about the existence of Gods X, Y and Z, the Wager cannot provide a compelling pragmatic reason for (acting so as to facilitate) believing in God X, as she might then risk infinite misery from God Y and/or God Z – one of which could even be a 'Perverse God' who

rewards those who believe only what can be established on the basis of objectively compelling evidence and argument... [and] whose wrath is reserved for those who ... would bring the canons of practical rationality to bear on epistemological matters. (Rosen, Undated)

It seems the live hypothesis idea can be incorporated into Wager scenarios which evade this 'many Gods' objection. James claims the Wager would only convince someone with a single prior living option concerning the God in question. If the

⁴⁶ This objection dates back to Diderot at least: 'An Imam could reason just as well this way' (Diderot, 1875-1877:167), (Hacking, 1972:188); which James also alludes to: 'As well might the Mahdi write to us ...' (1896, 2000:201).

individual has no other living option concerning another God then, for that individual, the 'many Gods' objection falls away, as there would be no need to choose between Gods. But even if the 'many Gods' objection is evaded the 'Wager God' needs other key attributes. It must be able and willing, indeed guaranteed, to operate as the Wager prescribes.

James himself doubts whether a God guaranteed to reward a believer only convinced of its existence by 'mechanical calculation' would be a God someone encountering the Wager might have a prior live hypothesis about: 'if we were ourselves in the place of the Deity, we should probably take particular pleasure in cutting off believers of this pattern from their infinite reward' (James, 1896, 2000: 201). To take the Wager seriously as a wager it seems reasonable to envisage someone accepting it willingly without fearing her willingness would disqualify her from winning, even if accepting it only meant deliberately acting so as to make receiving a God's grace more likely.

But perhaps this misrepresents Pascal's intention. We typically accept a wager voluntarily. Pascal however describes it as something we are engaged in anyway:

*...but you must wager. There is no choice, you are already committed.
...You would have to play (since you must necessarily play)... (Pascal, 1966:150-1)*

Hacking agrees:

*When reason cannot answer, a sensible man can say that he will not play the game. **But in our case, by the mere fact of living, we are engaged in play.** We either believe in God, or we do not. [Emphasis added.] (1972:188)*

James also interprets Pascal as seeing the aspirant's quandary not as *being*, but as being *like*, a gambler's:

*[Pascal is] reasoning **as if** our concern with truth **resembled** our concern with the stakes in a game of chance [Emphases added.] (1896, 2000: 201)*

Ordinarily you choose to accept a gamble or not, so it would not be a Jamesian forced option. But if the Wager is something we are engaged in anyway as well as something we could consciously and voluntarily engage in, then perhaps it is not introducing us to previously unimagined choices but explicating the risks we already undertake by choosing to live one way rather than another.

A paid-up Jansenist invented the Wager however, not an agnostic or atheist speculating about belief. There is a tension between the idea of a wager we consciously and voluntarily engage in and one we are 'already committed' to 'by the mere fact of living' because we have 'no choice'. The former could explain the Wager's exceptionalism: why belief in God, why God, why this God? Pascal constructed a wager around a specific deity whose existence he already accepted. It must be a specific wager, because we cannot consciously and voluntarily entertain an infinity of wagers. But if we, most likely with none of Pascal's background, are 'already committed', having 'no choice', to a wager on the existence of a Jansenist God we may be unaware of, then we must be already committed 'by the mere fact of living' to an infinity of other wagers, like:

(4.4AW1) At some time in the future but before your death a letter will arrive measuring 8 inches by 5 inches written in English in blue ink on yellow paper promising you eternal life. To accept the wager you must send a postal order for £10 by first class post to 23 Railway Cuttings, East Cheam. This must be done in advance, before the letter arrives.

(4.4AW2) As (4.4AW1), but the letter measures 7.9 inches by 5 inches.

(4.4AW3) As (4.4AW1), but the payment must be £11.

(4.4AW4...n) ...and so on, each with minor changes to one or more parameters.

If we 'must necessarily play' Pascal's concocted Wager about a supernatural character familiar around the European Counter-Reformation then we are also 'already committed' by the 'mere fact of living' to wagers (4.4AW1...n) and

countless others. Anyone could be seen as having chosen to live either in acceptance of (4.4AW1) or in explicit or implicit rejection of (4.4AW1), with ignorance counting as implicit rejection. Ditto for (4.4AW2), (4.4AW3) etc. The lack of anything 'moral' or 'spiritual' about accepting, rejecting or ignoring these wagers is irrelevant. Pascal uses moral language to describe the Wager's behavioural outcome: one who accepts it

will be faithful, honest, humble, grateful, full of good works, a sincere, true friend.... It is true you will not enjoy noxious pleasures, glory and good living, but will you not have others? (1966:153)

But this content follows from Pascal's conception of his Wager God. There is nothing necessary about this aspect to any particular wager. If we 'must necessarily play' Pascal's Wager then we must also 'necessarily play' (4.4AW1), (4.4AW2) etc. But we would not have live hypotheses in respect of many of them.

So enlisting James's live hypothesis to evade the 'many Gods' objection jars with the idea of a Wager we 'must necessarily play'. Perhaps we must play if we have the appropriate living option. So if we take the description of the Wager seriously we may find we are forced to play. But there is no necessity about taking the description of the Wager seriously and having the appropriate living option. So it is not true that we 'must wager', that there is 'no choice', that we are 'already committed' (Pascal, 1966:150) and 'engaged in play' (Hacking, 1972:188) by 'the mere fact of living'. Instead we can only see it as something we consciously and voluntarily engage in, or are at least committed to by virtue of a live hypothesis we might have but might not have had, in respect of a single specifically described deity. Pascal might think we all 'must necessarily play' but there is no reason to assume we share his ontological world-view.

This weakens, but does not eradicate, the threat to EP4. To evade the 'many Gods' objection the Wager needs something like James's 'live hypothesis', which means losing the universality and necessity Pascal claimed for it. So any threat to EP4 would need to be at the level of an individual explicitly or implicitly using Wager logic to decide on a course of action likely to result in an evidentially

unjustified belief in a God's existence. The challenge would then be to justify why EP4 would judge this decision illegitimate. To address this we need a potentially troublesome case. For clarification we will first consider a scenario which does not qualify.

In this first scenario individual A persuades herself to discount all other possibilities and entertain only two live hypotheses: (i) The Wager God exists; and (ii) The Wager God does not exist. She takes the Wager seriously, so understands she can only bet on one specific God. As this requirement cannot be forced on her, James claims that if she accepts the Wager, then the idea of a (single, prior) live hypothesis is needed to explain why she would find it compelling. She concludes though that the Wager God's necessary attributes include what to her seem counter-intuitive ones about rewarding a believer who only ended up believing by following a 'mechanical calculation', and punishing an individual who thought its attributes made it unworthy of worship and belief.⁴⁷ These attributes prevent her entertaining a live hypothesis about it. Hypothesis (i) is no longer live for her. She may have a live hypothesis about a God, but not about the Wager God.

She cannot vacillate between the Wager God and one she considers more worthy, because she can only wager on one God. She also wants her God to love and accept her if she just believed, not because she believed by accepting the terms of a wager. To commit to believing in the existence of a supremely worthy God, *outside* Pascal's Wager – perhaps from a conviction that doing so would achieve an otherwise unachievable moral benefit – might count as a Kierkegaardian leap of faith. But this would be believing in spite of the Wager, not because of it. In this scenario Pascal's Wager itself does not threaten EP4.

If we could see the Wager as something we are engaged in implicitly 'by the mere fact of living' whether we like it or not, then we might see its rationality from the outside. I might see someone else doing the rational thing by believing, because I see her betting on the existence of a specific God, which may or may not exist.

⁴⁷ '...Pascal's wager presupposes that God is a complete moral idiot' (Belzer, 2007:99).

But believer A cannot see that rationality herself, and cannot justify it in terms of probability calculations and decision theory, without seeing herself as actually engaged in the Wager, and therefore running the risk of seeing the Wager God as one unworthy of worship. In this scenario Pascal's Wager as a pragmatic argument risks corrupting the sincerity of belief it aims to promote.

Now for the potentially troublesome scenario. Individual B also persuades herself to discount every other possibility and consider only the two live hypotheses (i) and (ii).⁴⁸ But although B may be as aware as A of what the Wager God's attributes imply, she does not think the God unworthy of belief. She accepts the Wager, deciding a course of action likely to facilitate over-belief.

This may not be straightforward. Pascal advises someone intent on being 'cured of unbelief' not to go it alone but 'learn from those who were once bound like you' (1966:152). She is more likely to succeed by engaging with other believers in a shared God. Masses are communal and artefacts like holy water need a context of 'imitation and partisanship' and 'circumpressure of ... caste and set' (James, 1896, 2000:203). The shared God could be the Catholic God or Vishnu or Yahweh, not one abstractly concocted from Wager attributes. But we will assume we can navigate this. Perhaps she takes the shared Catholic God's existence to be a live hypothesis, takes the Wager's cost-benefit analysis to be consistent with her understanding of that God, and thinks coming to believe in this way is conceptually consistent with her faith since she believes it will open her up to that God's grace.

EP4's final verdict on B would depend on case details. Other things being equal EP4 would judge any voluntary action to facilitate over-belief morally wrong because of insufficient evidence, and therefore in a social context could encourage further gullibility⁴⁹ (and, in the case of religious over-belief, possibly cruel and repressive behaviour⁵⁰) in the believer and/or her community. B's potential prize cannot mitigate this, as the Wager's stake is a pragmatic, not a

⁴⁸ See also p159ff.

⁴⁹ See pp18–21: 'encourage gullibility' is here shorthand for behaviour which impedes the aims and benefits articulated in (1.1.1–13) and (1.2.1).

⁵⁰ See p136.

moral benefit.⁵¹ The winner maximises her personal utility (eternal life in paradise) rather than, say, net happiness in her community or universe. According to the Wager this is the individual's reason for accepting. The behaviour the individual commits to in accepting the Wager can have moral consequences: 'faithful, honest, humble, grateful, full of good works' (Pascal, 1966:153), which although strictly speaking outside the terms of the Wager, could be included in mitigation and then qualify as a potential 'conflicting and overriding moral imperative'. So if there were moral benefits these might exculpate B's over-belief, but if not EP4 would rule the over-belief morally wrong, other things being equal. None of this undermines EP4 as a proposed moral principle, or that an over-belief arrived at via Pascal's Wager represents any new threat to EP4.

Could there however be a threat to EP4 relating to B's putative right to 'choose to believe' in accordance with Pascal's Wager to get the chance of eternal life in paradise? This would resemble a right which, under certain conditions, James asserts against CP: to

refuse obedience to the scientist's command to imitate his kind of option, in a case where my own stake is important enough to give me the right to choose my own form of risk (1896, 2000: 215)

Consider again the case where the over-belief itself does not provide, or does not clearly provide, any moral benefit. (We can discount as unproblematic cases which EP4's 'conflicting and overriding moral imperative' caveat would exculpate.) The over-belief could not count as 'private',⁵² if only because of behavioural manifestations like masses and holy water. It must however be 'apparently harmless',⁵³ otherwise we could argue the potential negative consequences already identified, for over-belief in general⁵⁴ or religious over-belief in particular,⁵⁵ would override the individual's putative 'right to choose'. In **Chapter 3**⁵⁶ we suggested we cannot know how harmless an individual over-belief is until every

⁵¹ See p147.

⁵² See p122.

⁵³ p121.

⁵⁴ See pp14–17.

⁵⁵ See pp131–140.

⁵⁶ p122.

eventual consequence has materialised, hence Clifford's 'however seemingly trivial the belief ... we have no choice but to extend our judgment to all cases of belief whatever' (1877, 1879:182). We allowed that in an individual case we may have reason to think the over-belief would prove harmless in the long run, so a 'conflicting and overriding moral imperative' could legitimise it. Yet here the over-belief provides no moral benefit: does that rule out the conflicting and overriding moral imperative?

Not necessarily. We are asking if B has any right to win the chance of eternal paradise by accepting Pascal's Wager and choosing to believe. We are stipulating there are no moral benefits but at the same time that any other consequences are as innocuous as possible. So if she claims a right which EP4 infringes that suggests a conflicting and overriding moral imperative to protect that right. The over-belief does not have to have a moral benefit for this condition to apply. It could also apply to avoid a net moral cost.

In the context of Pascal's Wager and EP4 there would need to be a powerful moral argument in favour of honouring an individual's right to over-believe, in highly exceptional circumstances, to achieve a pragmatic benefit for herself rather than a socially contextualised moral benefit. This right could plausibly be regarded as a facet of a more general right to individual liberty, where no harm, or minimal harm, is likely to be incurred by a third party.⁵⁷ The over-belief would need to be free of any trace of 'pestilence' against which Clifford promotes CP, including potential repression⁵⁸ which religion can foster. But EP4 does not rule it out on principle, although it would rule out any blanket right to over-believe to achieve a pragmatic benefit for the believer rather than a socially contextualised moral benefit.

Our conclusion is therefore that we have not yet found a reason why Pascal's Wager should threaten EP4. Something like James's 'live hypothesis' seems necessary to protect the Wager from the 'many Gods' objection, but at the price of

⁵⁷ We will revisit this later (p195ff) when comparing EP4 with alternative potential principles in the context of religious belief.

⁵⁸ p136.

the universality and necessity Pascal claimed for it. Because of the role live hypotheses play in James's own account of what belief is permissible, we must now see if 'the will' can influence whether a hypothesis is live and, specifically, if we can voluntarily act so as to have the 'live hypothesis' we want to have.

This matters in our defence of EP4 since if we can will ourselves to have the live hypothesis we want, this could introduce circularity into James's argument for doxastic permissiveness.

Belief and the will

When James applies his framework to 'the actual psychology of human opinion' (1896, 2000:200) he rejects the idea of deciding to believe something purely by willing it; and, in at least some circumstances,

talk of believing by our volition ... is worse than silly, it is vile. (1896, 2000:201-202)

Yet he does not think we believe just by 'intellectual insight' or 'pure reason' (1896, 2000:203) either, particularly in a context like religious belief.

According to James:

(4.5.1) My will can only choose a hypothesis if it is already live for me.

Now if it was also true that

(4.5.2) I can make a hypothesis live for myself by an act of will,

then (4.5.1) would be trivial. I could always will to choose a hypothesis as long as I had already willed it into life. For (4.5.1) to be material, (4.5.2) must be false, indeed 'preposterous' (1896, 2000:200). Yet James also sees a role for 'our willing nature' in determining what are live and dead hypotheses for us:

*When I say 'willing nature,' I do not mean only such deliberate volitions as may have set up habits of belief that we cannot now escape from,—I mean all such factors of belief as fear and hope, prejudice and passion, imitation and partisanship, the circumpressure of our caste and set... **that make hypotheses possible or impossible for us, alive or dead.** [Emphasis added.] (1896, 2000:203)*

While some are 'deliberate volitions', others we have only partial or indirect voluntary control over. But they all count as non-rational sources, in contrast with the rational intellect. We can have some control over what we 'fear and hope'; we can nurture or dissolve our prejudices, and partially control whom or what we imitate or are partisan with, and the 'set' whose 'circumpressure' may influence us. But none of this implies we can voluntarily make a hypothesis live or dead directly, although we may choose to live in ways which make it more or less likely that some hypotheses will be live for us and others dead.

James's 'willing nature' is not therefore an exact synonym of 'conscious volition' or 'voluntary control'. But we can ask whether, from his pastry cart of 'influences', he is trying to have his cake and eat it:⁵⁹ our willing nature appears to be a faculty where we both do and do not exercise conscious and deliberate control. Willing is typically a matter of conscious and deliberate volition. But fear can reduce effective volition. Without fear I could freely choose P or Q. But fearing P's potential consequences could make choosing P harder than choosing Q, so I cannot choose P as freely as Q. Something equivalent applies to at least some of his other 'influences'.

Nor are we slaves to fear, prejudice and passion. We can learn about what influences us and do something about it. James's list includes phenomena where will is both exercised and constrained or compromised, but all are aspects of our 'willing nature'.⁶⁰

⁵⁹ Gale (1999, 2007:4) thinks James wants 'to have it all'.

⁶⁰ Gale (1999, 2007:235) identifies an 'aporia' running 'throughout James's general account of belief. On the one hand, his sentiment of rationality doctrine stresses that our beliefs are determined by our emotions and passions, which renders them non-intentional. On the other hand, his Promethean will-

We can also ask how long does (can, must) a hypothesis stay 'possible or impossible for us, alive or dead', following the influence of one or more 'factors of belief'. The next moment a new factor (a deliberate volition; or a fear, hope, or passion) could change the settings of what is 'possible or impossible..., alive or dead' for me. So even if we accept James's framework, we can question the ultimate significance of his distinction between live and dead hypotheses. Saying it is 'only our already dead hypotheses that our willing nature is unable to bring to life again' (1896, 2000:203) suggests that once a hypothesis is dead it can never be live again; so permanence seems part of the description and significance of the distinction between 'live' and 'dead'. But if a new 'factor' can change the settings of what is live or dead for me, then an aspect of our 'willing nature' can bring an already dead hypothesis to life again, so the setting is not permanent.

This matters for his treatment of Pascal's Wager. If the Wager option is 'act so you will probably come to believe in God' versus 'do not so act',⁶¹ then James's demand for a prior live hypothesis could fall away. Choosing 'masses and holy water' over getting 'yoked together with unbelievers'⁶² could count as 'action[s] of our willing nature' (particularly 'imitation and partisanship' or 'circumpressure of our caste and set') which 'make hypotheses possible or impossible for us, alive or dead'. They would then be legitimate enablers of live hypotheses in their own right, needing no 'pre-existing tendency to believe' (1896, 2000:201).

The overall line of argument in the Wager context displays either circularity or redundancy. We will take what James says about live hypotheses at face value, and also the Hacking/Franklin line (which is probably Pascal's and also James's line) that 'what the Wager motivates is not belief, but action' (Franklin, 1998:110). James says the individual needs a single prior live hypothesis about a God's existence to be swayed by the Wager. But if she is swayed, the outcome will be at most a decision to act so as to facilitate acquiring, and perhaps reinforcing but without confirming, that same live hypothesis. Hence perhaps why James calls the

to-believe doctrine requires that we be able to choose our beliefs at will, something we accomplish by making the intentional effort to concentrate our attention in a certain way...'.
⁶¹ p145ff.

⁶² 2 Corinthians 6:14 (New International Version).

Wager 'a last desperate snatch at a weapon against the hardness of the unbelieving heart' (1896, 2000:201).

More generally however the suggestion that we can choose to do something to influence, however indirectly, the live hypotheses we end up with exposes the 'live hypothesis' idea to a more critical light.

We might associate duties of responsible inquiry⁶³ more with Clifford, whose opening chapter is 'The Duty of Inquiry' (1877, 1879:177). But James also calls us 'would-be knowers'⁶⁴ (1896, 2000:209). If we are seekers after truth, why should our possession of H_1 as a live hypothesis give us the right to believe H_1 , as long as the other W2B conditions are met? What is special about H_1 other than our entertaining it? Alternatively, if we assume the right to have (or keep) H_1 as a live hypothesis, should we not have some rationale justifying why H_1 is live for us but H_{2-n} are not live for us, over and above the contingent fact that H_1 happens to be live and H_{2-n} happen not to be? Although a plausible duty of responsible inquiry cannot insist we investigate everything all the time (implying an infinite and therefore impossible task), it might entail that if our objective is to 'find out whether X' then we should make a reasonable effort to consider potential truths about X, and not rest content with the ones we currently happen to have in our heads.

This in turn suggests James's 'live hypothesis' may give Pascal's Wager a less effective defence than originally appeared: if you insist on sticking with the God you happen to have in your head or it does not occur to you that there could be more than one possible God, can you claim to be addressing the question seriously?

Individual B⁶⁵ was the nearest to a Wager scenario threatening EP4. Whether or not she qualifies as the 'responsible inquirer' EP4 presupposes might turn on how she persuades herself to 'discount every other possibility' leaving herself with one living option.

⁶³ May also be of relevance to 'moral encroachment', which is outside our scope: see p295.

⁶⁴ See p163ff.

⁶⁵ p153ff.

We raise this doubt not only in connection with James's reinterpretation of Pascal's Wager but because of its significance when we discuss whether his own 'religious hypothesis' represents a plausible threat to a principle like EP4.⁶⁶

Our passional nature

We concluded above⁶⁷ that the Wager, with or without James's help, seems no threat to EP4. But outside the Wager context, and discounting for now the doubt we have just raised, James's 'live hypothesis' seems straightforward enough for the job he recruited it for.⁶⁸ Before getting into his own treatment of religious belief, though, we should review a few remaining aspects of his framework. Following his own sequence we now consider his more normative remarks about how we decide what to believe.

The psychology of belief, says James, is

evidently far from simple; and pure insight and logic, whatever they might do ideally, are not the only things that really do produce our creeds. (1896, 2000:205)

He asks if this 'mixed-up state of affairs' is a 'normal element in making up our minds' what to believe or 'simply reprehensible and pathological'. It could of course be both. 'Normal' has both descriptive and normative connotations, so he could be saying either:

- (i) That is how we are, so we should resign ourselves to its wrongness and treat it as 'normal' [normal but wrong];

or:

⁶⁶ See p171.

⁶⁷ p155.

⁶⁸ It also resonates with the Bayesian notion of 'personal probability' to be discussed later: p213ff. It may be that only descriptive live hypotheses can qualify as Bayesian personal probabilities, and not prescriptive (normative) live hypotheses.

- (ii) If that is how we are then it must be 'normal' and cannot be 'reprehensible and pathological' [normal therefore not wrong.]

We will assume his intention is more (ii) than (i), as this is where he states his central thesis:

W2B *Our passionate nature not only lawfully may, but must, decide an option between propositions, whenever it is a genuine option that cannot by its nature be decided on intellectual grounds; for to say, under such circumstances, "Do not decide, but leave the question open," is itself a passionate decision,—just like deciding yes or no,—and is attended with the same risk of losing the truth. (1896, 2000:205)⁶⁹*

Empiricism, absolutism and objective evidence

James assumes we are eschewing 'systematic philosophical scepticism' (1896, 2000:205)⁷⁰ and instead holding as an article of 'faith' that 'truth exists, and that our minds can find it'. We can do this in either an 'absolutist' or an 'empiricist' way. An absolutist says 'we not only can attain to knowing truth, but we can *know when* we have attained to knowing it'; whereas for an empiricist 'although we may attain it, we cannot infallibly know when'.

James sees a distinguishing feature of absolutism in the 'absolutist conviction in a doctrine ... of objective evidence' (1896, 2000:206). Superficially this reads like a requirement of CP or EP4 but James has a stronger doctrine in mind, in that 'objective evidence' necessarily guarantees ('infallibly' supports) certainty ('bottom-certitude'):

*If ... I am **unable to doubt** that I now exist before you, that two is less than three, or that if all men are mortal then I am mortal too, it is because these*

⁶⁹ Repeated from p140.

⁷⁰ This does not imply that an individual faced with a specific option and choosing to 'not decide, but leave the question open' is thereby embracing 'systematic philosophical scepticism'.

things illumine my intellect irresistibly. The final ground of this objective evidence possessed by certain propositions is the [‘agreement of our intellect with the thing known’]. *The **certitude** it brings involves an* [‘aptitude for extorting a certain assent from our intellect’] *on the part of the truth envisaged, and on the side of the subject a* [‘quiet rest in knowledge’], *when once the object is mentally received, that leaves **no possibility of doubt** behind;...* [Emphases added and translations substituted for original Latin.] (1896, 2000:206)

This sense of ‘objective’ seems close to ‘conclusive’. To a modern ear, conclusive evidence may need to be objective, but not vice versa. I could claim subjective evidence that it will rain tomorrow, by feeling it in my bones or having a private visitation from a weather god. Or I could have objective evidence in the form of meteorological data about an approaching cold front. But although I believe it will rain tomorrow based on objective rather than subjective evidence, I am not ‘unable to doubt’ my belief. So we should rather speak of conclusive evidence, a doctrine James clearly contrasts with his own fallibilism:

I am ... myself a complete empiricist so far as my theory of human knowledge goes. I live ... by the practical faith that we must go on experiencing and thinking over our experience, for only thus can our opinions grow more true; but to hold any one of them—I absolutely do not care which—as if it never could be reinterpretable or corrigible, I believe to be a tremendously mistaken attitude... (1896, 2000:207)

This could be Clifford though, reminding us that ‘no greater service can be rendered to science than the purification of accepted results from the errors which may have crept into them’ (1877, 1879:197), and that ‘our bounden duty’ is that of ‘testing and inquiring into things’ (1877, 1879:199) and ‘purifying and enlarging ... to the utmost of our power’ (1877, 1879:205) the ‘precious deposit’ inherited from our forebears. James however counts Clifford among the ‘absolutists’:

When the Cliffords tell us how sinful it is to be Christians on such ‘insufficient evidence,’ insufficiency is really the last thing they have in mind. For them

the evidence is absolutely sufficient, only it makes the other way. (1896, 2000:206)

There seems little justification for equating Clifford's 'sufficient evidence' with James's 'objective', hence conclusive, evidence. James could however accuse Clifford of failing to define 'sufficient evidence'. As Hall summarises:

With respect to evidential sufficiency, nowhere does Clifford specify ... an infallible mark by which one could know that it has been reached; he only assumes that there is one. For James ... the quest for that criterion is a fool's errand. (2011:82)

Chapter 5 will address this challenge on behalf of EP4. If the certainty of conclusive evidence must be our criterion of evidential sufficiency then EP4 would outlaw the vast majority of beliefs we successfully live by, however 'reinterpretable' and 'corrigible' they may ultimately be (and however aware of this we may be).

Knowing the truth and avoiding error

James's last 'preliminary' is his 'two ways of looking at our duty in the matter of opinion', two 'first and great commandments as would-be knowers': '*We must know the truth; and we must avoid error*' (1896, 2000:209). These are 'not two ways of stating an identical commandment' but 'two separable... [and] materially different laws'. Hall (2011:82) calls them the 'counsel of hope' and the 'counsel of fear' respectively, and James caricatures Clifford in terms of the latter:

Believe nothing, he tells us, keep your mind in suspense forever, rather than by closing it on insufficient evidence incur the awful risk of believing lies.

... [This] merely shows his own preponderant private horror of becoming a dupe (1896, 2000:209-210).

James's distinction is understood to be about inductive risk. The more liberal your policy for accepting a hypothesis on the basis of evidence the more likely you are to accept truths but the more you risk accepting falsehoods. The more stringent your policy the more likely you are to avoid falsehoods but the more you risk rejecting truths.

For example a 99% reliable medical test comes back positive. If you believe on this basis, and adopt a general strategy of believing the results of tests which are $\geq 99\%$ reliable, you will gain many true beliefs but probably some false beliefs too. If you only accept the results of 100% reliable tests you will avoid false beliefs but lose many truths.

Given this interpretation we would need EP4 to encourage a reasonable balance of inductive (and, more generally, doxastic and/or epistemic) risk rather than insisting our only priority is to 'avoid error' as this could mean disbelieving many truths. Indeed EP4 explicitly outlaws believing 'irresponsibly' on insufficient evidence, so would accommodate believing 'responsibly' on the basis of evidence which is less than 'conclusive'. However a more satisfactory defence of EP4 needs that adequate criterion of 'sufficient evidence' which Hall calls for. **Chapter 5** will address inductive risk in more detail, but in the meantime it is not difficult to see where James and Clifford diverge.

In the context of much scientific and everyday belief they seem fairly aligned. Clifford talks of the

many cases in which it is our duty to act upon probabilities, although the evidence is not such as to justify present belief[.] (1877, 1879:189)⁷¹

He seems to withhold the commitment of belief proper in such cases. James also allows that in both 'scientific questions' and 'human affairs in general' scenarios exist where 'the option between losing truth and gaining it is not momentous' and

⁷¹ See p112.

we can throw the chance of gaining truth away, and... save ourselves from any chance of believing falsehood, by not making up our minds at all till objective evidence has come. ... Throughout the breadth of physical nature facts are what they are quite independently of us[.] ... The questions here are always trivial options, the hypotheses are hardly living (at any rate not living for us spectators), the choice between believing truth or falsehood is seldom forced. The attitude of sceptical balance is therefore the absolutely wise one if we would escape mistakes. (1896, 2000:210-211)

James and Clifford differ though in contexts where James sees belief as legitimate despite 'insufficient evidence', because we 'may be interested at least as much in positively gaining truth as in merely escaping dupery' and therefore cannot 'wait with impunity' for 'coercive evidence'.

James identifies three such 'forced options in our speculative questions' (1896, 2000:212). We have already covered 'moral beliefs'⁷² and 'questions concerning personal relations'.⁷³ The third is 'the religious hypothesis' (1896, 2000:212-214), the subject of this chapter.

We have discussed why James thinks Pascal's Wager is vulnerable without his 'live hypothesis' idea. We then touched on further elements of James's belief framework which will feature in his own defence of religious belief: the role of our 'passional nature' in belief formation; epistemic fallibility; and the tension between 'know the truth' and 'avoid error'. We are now ready to see if his defence threatens EP4.

4.4 James's case for religious belief

James's defence of religious belief has both a specific and a general aspect. He articulates his own 'religious hypothesis', but also offers his W2B thesis⁷⁴ as a more general 'defence of our right to adopt a believing attitude in religious matters'

⁷² See p57ff.

⁷³ See p110ff.

⁷⁴ See p140.

(1896, 2000:198). W2B therefore applies not only to his own religious hypothesis but to any hypothesis forming part of a 'genuine option that cannot by its nature be decided on intellectual grounds'. His own religious hypothesis is one example, but someone else's, different, religious hypothesis could be another.

James's religious hypothesis

James's religious hypothesis ('JRH') may not encapsulate the common essence of all religious beliefs, assuming that exists. But it shares an important structural feature with most. While science 'says things are' and morality 'says some things are better than other things', religion says

the best things are the more eternal things, the overlapping things, the things in the universe that throw the last stone, so to speak, and say the final word. "Perfection is eternal" ...seems a good way of putting this first affirmation of religion, an affirmation which obviously cannot yet be verified scientifically at all.

The second affirmation of religion is that we are better off even now if we believe her first affirmation to be true. (1896, 2000:214-5)

Also for 'most of us',

[t]he more perfect and more eternal aspect of the universe is represented in our religions as having personal form. The universe is no longer a mere It to us, but a Thou, if we are religious; and any relation that may be possible from person to person might be possible here. (1896, 2000:216)

These three components combine descriptive beliefs about how things are with prescriptive beliefs about how things should be. The two kinds of belief are also connected, similar to our (4.1.1–3).⁷⁵ It is a necessary structural feature (of at least this kind – our kind and James's kind) of religious belief that they are connected.

⁷⁵ See p131.

Descriptive content includes:

- (4.6.1.1) The universe has 'more eternal' and 'overlapping' things which 'throw the last stone' and 'say the final word'.⁷⁶
- (4.6.1.2) There is a link between what is 'more eternal' and what is 'more perfect'.
- (4.6.1.3) 'The more perfect and more eternal aspect of the universe' has at least some features of a person.

Prescriptive content includes:

- (4.6.2.1) We ought to believe (4.6.1.1) and (4.6.1.2).
- (4.6.2.2) We ought to believe and behave in accordance with (4.6.1.3), because 'by obstinately believing that there are gods' ('more perfect and more eternal' aspects of the universe with at least some features of a person) 'we are doing the universe the deepest service we can' (1896, 2000:216).
- (4.6.2.3) We should behave in accordance with the 'religious hypothesis': there are kinds of 'action required or inspired by the religious hypothesis' which are different from (and, by implication, potentially better than) what is 'dictated by the naturalistic hypothesis' (1896, 2000:217 footnote).

Religious belief as a genuine option

In view of the inseparably compound nature of religious belief, exemplified by 4.6.1.1–4.6.2.3 above (and by 4.1.1–3⁷⁷), it is perhaps unsurprising that James's

⁷⁶ Hall reads 'the best things are the more eternal things' as 'good will out' (2011:84): *sub specie aeternitatis* goodness will triumph.

⁷⁷ p131ff.

'options' include choices both between potential beliefs and between potential actions.⁷⁸ As one of the 'great cosmical matters' (1896, 2000:214), religious belief combines decisions about what to believe about how things are with decisions about what to do, since a belief about what should be is, in a personal and moral context, tantamount to a decision about what is to be done, and therefore what to do.⁷⁹

This compound nature could also be why James insists the religious belief option is 'genuine'. For someone faced with the decision it is a **living** option. James would call it **forced**, because if a religious hypothesis is a candidate truth, the dilemma is 'Either accept this truth or go without it' (1896, 2000:199). Someone saying 'yes' could see the option as forced, for example between having and not having God in her life; and therefore see the practical (emotional, spiritual, psychological, existential, behavioural etc) consequences of saying 'no' equivalent to those of not deciding. But someone saying 'no' or 'don't know' or nothing might instead have more reason to see the option in terms of, say, morality or integrity, so might see 'no' and 'don't know' as distinct both from each other and from no response at all. An atheist rejecting a God hypothesis because of its potential moral harm could still be plagued by doubt, and envy one agnostic who discounts it on grounds of insufficient evidence and another who cares too little either way.

This could undermine James's claim that religious belief arising from a 'genuine option' is permissible. Can I be faced with a forced option just because I see myself faced with a forced option? For James,

if I say, "Either accept this truth or go without it," I put on you a forced option, for there is no standing place outside of the alternative. Every dilemma based on a complete logical disjunction, with no possibility of not choosing, is an option of this forced kind. (1896, 2000:199)

⁷⁸ See p143.

⁷⁹ Cf discussion whether Pascal's Wager is about believing something is true or deciding to act in a certain way p146ff.

'Either accept this truth or go without it' certainly reads like a forced option. We could question though how significant that reading is in respect of a purely descriptive 'truth'. James does not elucidate what 'Either accept this truth or go without it' means in practice, as it is not immediately obvious.

Compare 'Either accept this cake or go without it'. Here you implicitly agree that 'this cake' refers to a specific cake which you either accept or not. But you cannot refer to a specific hypothesis as 'this truth' without implicitly accepting it as true, because if you did not accept it as true you would not recognise it as something which could be described as a 'truth'. You cannot simultaneously recognise hypothesis H as 'this truth' and 'go without' H as a 'truth'.

This suggests we should interpret James's forced option as 'Either accept this as truth or go without it'. But being confronted with this does not seem wildly different in practice from being confronted with 'Either call my theory true or call it false' (1896, 2000:199), which for James is avoidable, because one could also express no opinion either way. There is also more than one alternative to 'accept H as truth', because there is more than one way to 'go without H'. You could think H is false or have no opinion about H.

This suggests that an individual happy to entertain 'Either accept H as truth or go without H' literally as the forced option it is worded as must also be happy to discount any differences between alternative ways of 'going without H'. But that seems odd in respect of a purely descriptive hypothesis. Why would someone agree to, say, 'Either accept the hypothesis that gold's atomic number is 79 as truth or go without this hypothesis' if that meant discounting any difference between, say, believing gold has an atomic number but not knowing what it is; sincerely believing gold's atomic number is not 79 but 75; not knowing what an atomic number is; not knowing what gold is; etc? A Jamesian explanation could be because the option is also momentous since, like the decision to join Nansen's expedition it would involve a once-in-a-lifetime opportunity with life-changing consequences. But the Nansen option is about deciding what to do, not what to believe.

On hearing the speaker's words the addressee may accept it as a forced option: either believe gold's atomic number is 79 or take any other course of action (or inaction) logically incompatible with believing gold's atomic number is 79. But the addressee could instead see herself faced with a choice between, say, (i) believing gold's atomic number is 79; (ii) believing it is not 79 but another value; (iii) believing it may be 79 but needing further proof it actually is 79; (iv) remaining ignorant as to what an atomic number is, and/or what gold is; ... and so on.

In the case of a purely descriptive hypothesis about gold's atomic number the decision seems in practice more like one we 'may decline to offer any judgment' about, like 'Either call my theory true or call it false'. There seems nothing in the hypothesis itself to force an individual to see it as part of a complete logical disjunction, and therefore as part of a forced option.

Compare this with the decision whether or not to have God in your life. Assuming you only have one life, then that fact would make it a forced option: have God in your life or do not have God in your life. This is different from a decision purely whether to believe a God exists, which in itself does not have to be part of a forced option. But an individual could choose to see it as equivalent to, say, a decision about having or not having God in her life, which would present itself as a forced option to anyone with only one life.

So it is perhaps easier to understand why an individual might choose to see a hypothesis about the existence of a God as part of a forced option. It is easier to see how a decision with potential behavioural and/or existential implications can be forced for an individual, for example in relation to religion, where we are 'supposed to gain ... by our belief, and to lose by our non-belief, a certain vital good' (1896, 2000:215). But the challenge is to understand how religion can be 'a forced option, so far as that good goes' which can be 'put on' one person by another. This is so as to understand whether and, if so, how something which appears a forced option to one individual must appear forced to anyone else:

We cannot escape the issue by remaining sceptical and waiting for more light, because, although we do avoid error in that way if religion be untrue,

we lose the good, if it be true, just as certainly as if we positively chose to disbelieve. (1896, 2000:215)

Considering this quotation in isolation, the first person plural suggests James sees a forced option as applicable to anyone. But James has just said he is only addressing those for whom '*the religious hypothesis in both its branches*' is a live hypothesis:

If we are to discuss the question at all, it must involve a living option. If for any of you religion be a hypothesis that cannot, by any living possibility be true, then you need go no farther. I speak to the 'saving remnant' alone.

It is unclear how someone can 'put on [me] a forced option' if I am outside that 'saving remnant', do not see myself faced with a forced option, but instead see myself free to 'escape the issue'. A checkmated chess player has the forced option between conceding defeat or losing her king. But she was not forced to play.

If an option is only forced for me if I accept it as a forced option, while someone else can avoid it by not accepting it as a forced option, this would seem to reduce the significance of its 'forced' nature, and therefore weaken the claim that belief is permissible in respect of an option which is forced (among other conditions), as the following discussion will try to explain.

To say an option may be (or appear) forced for me while not everyone need see themselves faced with that same forced option does not imply I am consciously and/or voluntarily choosing to see myself faced with a forced option. In James's framework my options derive from the hypotheses which are live for me, which I have not necessarily chosen to have. But we have established⁸⁰ that whether or not a hypothesis is live for me can indirectly result from an action of my 'willing nature'. As a member of James's 'saving remnant' (an 'insider') I may see myself faced with a forced option between (i) believing a religious hypothesis and thereby

⁸⁰ p158.

gaining a 'vital good' if religion is true, or (ii) potentially losing that vital good, again if religion is true, by either (iia) 'remaining sceptical and waiting for more light' or (iib) 'positively [choosing] to disbelieve'. Someone else however, even another 'insider', could have that same religious hypothesis, or another one, as a live hypothesis, but not see herself faced with a forced option – because she has a different mix of live and dead hypotheses from the ones I have.

So, even as an insider, I do have another alternative, which I am free to choose whether or not to take. That is (iii) to behave in ways which might directly or indirectly facilitate my having different live hypotheses, and therefore different options, from the ones I currently have or think I have.⁸¹

James could of course argue that (iii) is actually (iic): another variant of his scepticism/positive disbelief route. But he cannot have it both ways, which he seems to be trying to do by holding these three positions simultaneously:

- (4.7.1) basing an agent's right to let her 'passional nature' decide what to believe in respect of a 'genuine option that cannot by its nature be decided on intellectual grounds' (1896, 2000:205) on her individual and contingent mix of live and dead hypotheses and her individual and contingent evaluation of what options are forced and momentous for her;
- (4.7.2) insisting that that same 'passional nature' is at least a partial determinant of what these genuine options are; and
- (4.7.3) describing our 'actual psychology' such that we can voluntarily act in ways which might directly or indirectly influence that same 'passional nature'.

⁸¹ We are not claiming there are no forced options which are forced for everyone, just that the option to believe or disbelieve a religious hypothesis does not seem to be one of them.

We may or may not be able to change our 'passional nature' on any particular occasion, but this does not exempt us from trying, if we have reason (including 'intellectual grounds') for thinking this a wise course of action.

James's W2B doctrine on the other hand would let an individual believe an evidentially unsupported hypothesis as long as, in that particular moment, (i) it was live for her; (ii) its negation was also live and evidentially unsupported; and (iii) it appeared to her as part of a forced and momentous option. In its unqualified state W2B imposes no duty to question whether one's current set of live and dead hypotheses are sound, and whether any of the decisions one currently sees as forced options can be recast as avoidable.

A critic could complain that we are misinterpreting James, and that his point about forced options is purely about inductive risk: either accept this claim and the potential benefits and harms that might come with it or do not and lose those benefits or harms. Consider a secular hypothesis like:

(4.8.1) Vaccine X prevents disease Y.

The practical effects of not believing (4.8.1) either way might be the same as believing its negation.⁸² But here again what generates the forced option is a fact about the world in relation to a potential decision. A public health officer cannot both administer vaccine X and not administer vaccine X, regardless of how her eventual decision got made. She must decide what to do. But this in itself does not compel her to see descriptive hypothesis (4.8.1) as part of a forced option, as she could continue not to believe either way (perhaps her understanding of the clinical trials to date leads her to conclude its effectiveness is not 100% proven) even though she knows she must make a decision under risk and decides to administer the vaccine. Her 'passional nature' may be instrumental in forcing her to a decision about what to do before it is too late, but her professionalism may also constrain her to submit her mix of relevant live and dead medical hypotheses to regular evaluation.

⁸² There is a parallel here with the choice between prioritising the imperative to 'know the truth' and prioritising the imperative to 'avoid error' (James, 1896, 2000:209): see pp163–165.

The key point is that the officer may or may not see descriptive hypothesis (4.8.1) as part of a forced option, regardless of the fact that

(4.8.2) Administer vaccine X or do not administer vaccine X

cannot but be a forced option. In practical terms the truly unavoidable decision under risk is: either carry out this action and realise the potential benefits and harms that might come with it or do not and lose those benefits or harms.⁸³

The officer's forced option (4.8.2) would also be what James calls momentous, as the religious belief dilemma is for an insider like himself:

If religion be true and the evidence for it be still insufficient, I do not wish, by putting [Clifford's] extinguisher upon my nature ... to forfeit my sole chance in life of getting upon the winning side (1896, 2000:215-6).

The option would be momentous by being potentially life-changing for the individual faced with it and, in the case of vaccine X, potentially for an entire population. In an individual with free will the option whether or not to believe a 'God hypothesis' could be momentous if, for her, it is inextricably linked to decisions about how to live the rest of her life, including the decision whether or not to live the rest of her life with the hypothesised God in her life. It would be forced because she has only one life. There is 'no standing place outside' the two alternatives 'live the rest of your life with God in your life' and 'live the rest of your life without God in your life'.

Again compare 'Either accept this cake or go without it'. This is forced because it is about what to do. There are not multiple ways of 'going without' a cake, unlike the multiple ways of 'going without' a descriptive hypothesis. But going without cake is trivial, not momentous.

⁸³ A somewhat different perspective on cases similar to these will be discussed in **Chapter 5** under the heading of 'pragmatic encroachment': see p240ff.

The challenge now is to understand how a purely descriptive 'God hypothesis' can be part of a forced and momentous option. If we focus purely on the descriptive content and ignore any implications belief in a God's existence might have for how the individual lives her life, an option involving a 'God hypothesis' would no more need to be forced than any other option involving descriptive hypotheses, if only because 'atheism' and 'agnosticism' represent two different and distinct alternatives to belief, broadly equivalent to the alternatives available to someone reluctant to 'call my theory true'. To ensure the relevant hypotheses are purely descriptive we must exclude any behavioural implications of 'living without God in your life', either while entertaining the possibility that God may exist or because you believe God does not exist and therefore cannot be in your life.

We have two distinctions in mind, separate but overlapping. One is between a descriptive hypothesis (that something is the case) and a prescriptive hypothesis (that something ought to be the case). The other is between a descriptive hypothesis and a decision (or potential decision) to do something. Although James introduces 'option' as a 'decision between two hypotheses', many of his examples are decisions between two potential actions (or between doing something and not doing it).⁸⁴

These two distinctions overlap in that a prescriptive belief (that something ought or ought not to be the case) is typically a belief that something is or is not to be done, so in the context of personal voluntary decision-making it is a decision to do something or not do it. (The 'ought' is not necessarily a moral ought.) This is particularly evident in a religious context, where a personal decision to adopt a particular religion is typically a decision to live a particular kind of life and behave in particular ways, because this is considered the right thing to do – either morally right or pragmatically right or a combination of the two. It is plausible that a decision whether or not to live a particular kind of life, especially with the holistic implications typically associated with religion, can appear both forced and momentous to someone faced with the option: forced because it is a mutually

⁸⁴ See p143.

exclusive either/or (you have only one life, which you either live this way or not); momentous because it can be life-changing – in this world and perhaps the next.

But it is harder to understand how a decision about a purely descriptive hypothesis can be either as forced or as momentous, even where the descriptive hypothesis is about whether such-and-such omniscient and omnipotent entity exists and is the creator of all reality. Indeed it is hard to imagine how a purely descriptive hypothesis about an omniscient and omnipotent creator of reality can in practice stay free from prescriptive accretions. If you think it is true, or you think it could be true, you might understandably start thinking about possible consequences for yourself and what, at the very least, you should do to preserve your well-being.

Descriptive content ('is') may not generate prescriptive content ('ought') all by itself. But descriptive content can account for why the prescriptive content arose, in the sense that if the descriptive content had not existed then the prescriptive content would not have arisen in response.

Take this pair of connected beliefs:

- (4.9.1) There is an eternal, omnipresent, omnipotent, omniscient and judgmental God responsible for creating and maintaining all there is, and which knows everything about you including your innermost thoughts.
- (4.9.2) Therefore you should do everything in your power to appease this God.

The 'Therefore' in (4.9.2) represents not logical entailment but an understandable pragmatic response. In this sense we can say (4.9.1) accounts for (4.9.2).

If a 'religious hypothesis' was only about what is the case – for example (4.9.1) with no trace of (4.9.2) – it would seem less likely that the insider would see it as part of a forced option, and perhaps also with less justification. A purely descriptive hypothesis might have no implications for the individual's present or

future life: the option whether or not to believe being more like 'Either call my [metaphysical] theory true or call it false', and therefore avoidable.

But for a believer it could be the practical consequences for action and orientation, arising from component hypotheses about what ought to be (what is to be done, and therefore what to do), which help make it part of a forced option. Her 'sole chance in life of getting upon the winning side' would depend both on what she believes and on what she does. The structured combination of beliefs about what is and what ought to be (and so how to live her life) would also help make it momentous. She only has one life, to be lived either running 'the risk of acting as if [her] passional need of taking the world religiously might be prophetic and right' (1896, 2000:216), or not.

James's W2B permissiveness seems then to have an intriguing potential consequence in the context of religious belief. Even for an insider a purely descriptive religious hypothesis does not have to appear as part of a 'genuine option'. This is because for the individual to see the hypothesis as part of a momentous and/or forced option there would need to be implications for how she lives her life, guided by prescriptive content about what she should do. It would then be that prescriptive content which the descriptive content is bound to which makes the resultant nexus part of a living, forced and momentous option, and therefore yield for James a legitimate over-belief. However, as previously argued,⁸⁵ in the context of much religious belief the descriptive component is what makes the prescriptive component binding. This may not qualify as viciously circular reasoning. But it is worth identifying a potential scenario which James's W2B doctrine seems to allow. To avoid misunderstanding, we will refer to both prescriptive content about what ought to be and decisions about what to do as 'nondescriptive' content.⁸⁶

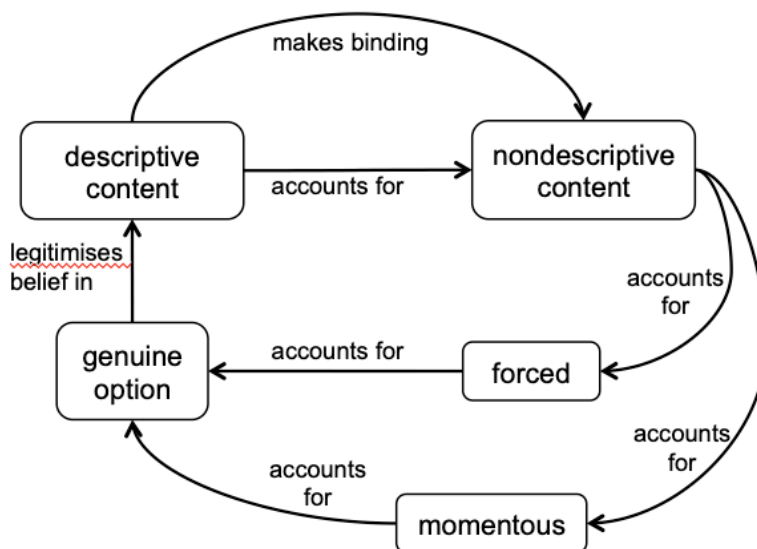
James derives the right to believe an evidentially unsupported live hypothesis from its being part of a forced and momentous option. An individual sees the decision whether or not to believe a religious hypothesis as a forced and momentous

⁸⁵ See p136.

⁸⁶ See p175.

option. The option is forced and momentous for her because of the nondescriptive content of the religious hypothesis, not because of its purely descriptive content. However in this scenario it was the descriptive content which accounted for the nondescriptive content: see (4.9.1–2) above. She then exercises the Jamesian right to believe the hypothesis, a right which covers both its descriptive and nondescriptive content. The descriptive content makes the nondescriptive content binding on her:⁸⁷ if the God does not exist then its commandments fall away. But her right to believe both the descriptive and the nondescriptive content derives from the option being forced and momentous, and it is only forced and momentous because of its nondescriptive content. If there was only descriptive content the option would not have been forced and momentous.

The diagram below attempts to summarise this:



It is possible of course that James himself might see nothing untoward in a sequence of steps like this.⁸⁸ But anyone who does may see a second major benefit in a principle like EP4, which would require an independent evidential justification for the descriptive content of a belief. Like the first⁸⁹ it follows from the

⁸⁷ See p136.

⁸⁸ He might indeed reject the distinction between the descriptive and nondescriptive content of a hypothesis.

⁸⁹ p137.

combined descriptive plus prescriptive content of religious belief – although in this case the direction of travel is from *ought* to *is* rather than *is* to *ought*.

Person to person

The conjoined is/ought nature of religious belief⁹⁰ is a concern with James's overall position: including the implications of W2B in the hands of someone without James's radical-empiricist leanings. And as already argued,⁹¹ the issue is the conjuring power of its descriptive component.

In the case of a religious hypothesis like James's JRH⁹² the descriptive component includes belief in the personal nature of divinity: 'The universe is no longer a mere *It* to us, but a *Thou*...' (1896, 2000:216).⁹³ *If* the God is a person, and therefore the relation between believer and God is between two persons, then we are in a realm where belief proper can slide into faith and trust. For example, by adapting (3.6.3)⁹⁴ we could get, as an address to a God who is assumed to exist:

(4.10.1) I believe that if I (succeed in making myself) believe that you will love me, this will tend to make you love me.

James maintains that in the context of such 'living options which the intellect ... cannot by itself resolve' we have a 'right to believe at our own risk any hypothesis that is live enough to tempt our will' (1896, 2000:217). In theory this could mean a yet more pre-emptively creative 'faith in a fact' (1896, 2000:213-4).⁹⁵

(4.10.2) I believe that if I (succeed in making myself) believe that you exist for me as a person, this will tend to make you exist for me as a person.

⁹⁰ See 4.2 Compound structure of religious belief p131.

⁹¹ See p136ff.

⁹² p166ff.

⁹³ See (4.6.1.3) p167.

⁹⁴ p114.

⁹⁵ See p110.

'Exist as a person' is convenient wording to distinguish it from, say, 'exist as a force of nature' or 'exist as an inanimate object'; and to accommodate James's characterisation that 'any relation that may be possible from person to person might be possible here' (1896, 2000:216) – which would include perceived duties and obligations.

The issue is that, depending on how we interpret the relevant parts of James's framework, I could end up with a legitimate belief that a God exists as a person. This is significant because, as argued in **Chapter 3**,⁹⁶ an interpersonal context can introduce additional 'faith in a fact' imperatives like those of trust and generosity which could exonerate instances of over-belief.

The context of (4.10.1) is that we are *assuming* the God exists and is a person ('*If the God is a person...*'). But if there are circumstances where we can legitimately believe the God exists and is a person, then under those circumstances *assumption* is unnecessary. Our working definition of belief is that 'our beliefs are what we regard as true' and 'what I believe is just how things are for me'. So if I (legitimately) believe that H, then H is (legitimately) what is true for me. James identifies circumstances where we can legitimately believe a hypothesis which is unsupported by evidence but forms part of a genuine option. Consider where H = 'God exists and is a person' and assume our working definition of belief. Now, if there are circumstances where I can legitimately believe that 'God exists and is a person' then in those circumstances 'God exists and is a person' will be a legitimate statement of how things are for me. So if I legitimately believe the God exists as a person, this means the God exists as a person for me.

If the God is a person for me, that means I would (should) relate to the God as one moral agent to another, and would have additional moral duties which I would not have were the God an impersonal object. As examples from **Chapter 3** illustrate ('Faith in a fact',⁹⁷ 'Smith and Jones',⁹⁸ 'Husband and wife'⁹⁹), these include additional doxastic duties, tending to favour belief over doubt. In the

⁹⁶ See **3.4 Personal relations and self-fulfilling beliefs** pp110–119.

⁹⁷ p110.

⁹⁸ p116.

⁹⁹ p117.

context of EP4 these additional duties would be subsumed under the 'conflicting and overriding moral imperative' condition.

Correctly applied however EP4 would disallow from inception any over-beliefs in the existence and personhood of God. This is therefore a third reason for adopting EP4 in respect of religious belief.

We must remember however that James speaks of our right to believe 'at our own risk' (1896, 2000:217). It is time to unpack this.

At our own risk

James sees religious belief as a purely private affair:

No one of us ought to issue vetoes to the other, nor should we bandy words of abuse. We ought, on the contrary, delicately and profoundly to respect one another's mental freedom: ... then only shall we live and let live, in speculative as well as in practical things. (1896, 2000:218)

This aligns with his discussion elsewhere of 'over-beliefs' being 'absolutely indispensable' to an 'individual's religion', and which

we should treat ... with tenderness and tolerance so long as they are not intolerant themselves. (1902, 1960:489-90)

If an individual's religious over-belief is purely private, only that individual bears the 'risk'. If James assumes this, then the *right to believe at our own risk* would equate to the *right to believe, which in the context of religious belief is purely at our own risk*.

If so it is not an assumption Clifford would share, with his 'credulous' societies reverting to 'savagery'. James could however stipulate that 'at our own risk' acknowledges and explicitly excludes scenarios where the effects of religious belief extend beyond the individual believer. In that case the quarrel between

Clifford and James would only apply in theory with no practical consequences.¹⁰⁰ At a practical level Clifford could allow the 'right to believe at our own risk' because he sees no hypotheses (i) which could be believed despite insufficient evidence, and (ii) where the potential net negative consequences would only impact the over-believer herself.

A Jamesian could counter that (ii), which would require evaluating the consequences of belief beyond the believer herself, itself represents a 'genuine option' involving how to balance the risk of consequential harm (to others) from the believer's credulity against the loss of consequential good (to others) from the potential believer's failure to believe. The Jamesian could say we have no unique, external criterion to decide for us, so we remain in the realm of choice, where '[o]ur *passional nature not only lawfully may, but must, decide*' (1896, 2000:205).

On behalf of EP4 we would agree we are in the moral domain – where moral positives and negatives net against each other. Clifford and EP4 see credulity as a bad thing, because of its corrupting influence. The total consequential harm from an individual act of over-belief may be major or minor, but will be non-zero – discounting any potential benefit of that same over-belief, which we must factor in. James sees a kind of benefit which can only accrue from religious over-belief:

Since belief is measured by action, he who forbids us to believe religion to be true, necessarily also forbids us to act as we should if we did believe it to be true. The whole defence of religious faith hinges upon action. If the action required or inspired by the religious hypothesis is in no way different from that dictated by the naturalistic hypothesis, then religious faith is a pure superfluity, better pruned away... (1896, 2000:217 footnote)

James seems to claim a necessary connection between religious belief and a particular kind of (desirable) behaviour. But do we have good reason to think such behaviour impossible without religious belief?

¹⁰⁰ We will pick this up later p198ff.

There are, say, Christians and Jews who do not believe in the literal existence of their God as an entity inside, outside or in any way related to the universe, but are Christians or Jews by subscribing to their understanding of the Christian or Jewish way of life, without over-beliefs. A subset of these will be morally good individuals because of conforming to the Christian or Jewish way of life. Other Christians and/or Jews will have over-beliefs about the literal existence of their God and be morally good because of those over-beliefs.

But James's case rests on this last group having no other option, because of something about religious belief. Given the actual and potential negative outcomes from over-belief, if James claims a general right to over-belief in a religious context because an 'action required or inspired by the religious hypothesis' could be 'different from that dictated by the naturalistic hypothesis' ('different' implying 'possibly superior as a result of that difference') he must identify something unique about religious over-belief which can make someone morally good in a way otherwise unachievable. We could grant this might be true for an individual: that for whatever reason she can only be as morally good as she is because of her over-beliefs. If so, then it would be morally right for her to over-believe. That is how moral decision-making goes. It would not undermine EP4. But to claim a general right to religious over-belief needs something universally unique about religious over-belief. This is what the footnote passage (1896, 2000:217) seems to claim. There is no quarrel if James is only comparing 'religious hypothesis' action with 'naturalistic hypothesis' action for an individual. Sometimes it is right to lie; sometimes even to kill.

James's position relies on there being a distinct category of beliefs where applying a principle like CP or EP4 would systematically result in a net loss of good. But all he seems to offer in support are (i) his belief in his own version of the religious hypothesis (JRH); (ii) his belief that his belief in JRH provides him with otherwise unachievable moral benefits; and (iii) his belief that his belief in JRH provides him access to an otherwise unachievable truth. These beliefs are no doubt sincerely held. But they are hardly enough to convince anyone outside James's own belief system that any of it justifies rejecting a principle like EP4.

James could retort that he is not offering objectively compelling reasons for everyone, just claiming more modestly that such belief is sometimes permissible. But then there is no conflict. EP4 'permits' individual instances of over-belief, where the over-belief is justifiably intended to achieve some other moral good which overrides the wrong of over-belief. If that is all James is claiming, then his permissiveness is 'compatible' with EP4 – as the moral rightness of lying to the Gestapo is 'compatible' with the otherwise general moral principle that lying is wrong.

Indeed, if the individual's circumstance was exactly that – that over-belief was the only way to achieve that net moral good – then she would morally *obliged* to over-believe, not just *permitted*. James is saying however that because the genuine option cannot be decided on intellectual grounds, the non-rational willing nature is obliged to choose, but permitted to choose either hypothesis. There is no obligation to choose one hypothesis over another. But in our extreme scenario where over-belief is permitted, it is permitted *despite* EP4, and is obligatory – because of one or more other moral principles.¹⁰¹

Otherwise we are merely acknowledging that the moral realm is one of choice and at times we face moral dilemmas because our principles give conflicting guidance. Do we have a uniquely correct way to resolve such dilemmas, for example rules about how to apply our principles together? We might posit a shared conception of the good, perhaps a utilitarian formula. But we may not agree this as a principle, or agree the formula. Despite our best intentions and ingenuity we may conclude that unresolvable moral dilemmas will always remain; and if so how do we decide what to do? What else can we employ (a Jamesian might ask) but 'our passionate and volitional nature' (1896, 2000:217)?

This rejoinder may be problematic for Clifford, but not for EP4. If we face a moral dilemma because the available options need us to choose between principles which in a specific context conflict, then any rationale we use to decide would at

¹⁰¹ We will have more to say about this later: p200.

most qualify as a prescriptive belief. It would be subject to CP, but not to EP4, which excludes prescriptive beliefs, and is itself a moral belief.

Although EP4 is a moral principle, it cannot resolve dilemmas arising from conflicting moral principles, including itself. A conflict between 'it is wrong to lie' and 'it is wrong to harm an innocent person' cannot be resolved using either principle: if it could there would be no dilemma.

But if it is a genuine moral dilemma (a Jamesian might then respond), this could be precisely where reason runs out and our 'passional nature' must take over. If so then there might be little distance between James and EP4, with the difference only one of emphasis: EP4, like Clifford, emphasising the potential social implications of encouraged and ingrained habits of over-belief; James assuming these can be ignored in the apparent privacy of religious belief.¹⁰²

Depending how we interpret James however there is one way of resolving a moral dilemma which he would apparently permit but which EP4 would disallow. This we must now consider.

Moral sovereignty

We have stressed that one reason religious belief is such a significant test case is its compound descriptive/prescriptive structure. A key aspect of this is the consequences for the believer's moral sovereignty. There may be restrictions over the kind of religious belief James's model would allow, depending on what 'at our own risk' means exactly.

James's W2B doctrine licenses the individual to believe any hypothesis forming part of a genuine option which cannot be decided on intellectual (including evidential) grounds. A hypothesis must be live for our will to choose it – and we have established¹⁰³ that whether or not a hypothesis is live for us can indirectly result from an 'action of our willing nature'.

¹⁰² See later p195ff where we discuss another potential rapprochement between James and EP4.

¹⁰³ pp156–160.

Take this 'living option':

Hypothesis (4.11.1) = the combination of (4.1.1) 'God exists' + (4.1.2) 'God is the source of morality, in that what God commands must be obeyed'

Hypothesis (4.11.2) = 'hypothesis (4.11.1) is false'

W2B would allow someone to have (4.11.1) plus (4.11.2) as a living option, even if (4.11.1) only became live for her as an indirect result of mixing with adherents who believed (4.11.1). W2B would let her believe either (4.11.1) or (4.11.2).

We will now expand (4.11.1) to include commandments to persecute non-believers and cede her moral sovereignty to God. The latter would mean a commandment she believes is from her God takes priority over any conflicting imperative from her own conscience.

But where does that leave James's injunction to tolerate each other's over-beliefs 'so long as they are not intolerant themselves'? If the tolerance imperative takes precedence, there are versions of the expanded (4.11.1) which James would not allow – particularly those including the imperative to cede one's moral sovereignty to God. So the individual's own moral position (were she to follow James's ethic of tolerance) would in this case trump an imperative from God. But this sits uneasily with James's 'The Moral Philosopher and the Moral Life' (1891, 2000):

*Our attitude towards concrete evils is entirely different in a world where we believe there are none but finite demanders, from what it is in one where we joyously face tragedy for an **infinite demander's** sake...*

*It ... is my final conclusion ...that the stable and systematic moral universe for which the ethical philosopher asks is fully possible only in a world where there is a **divine thinker with all-enveloping demands**. [Emphases added.] (1891, 2000:262)*

How would James respond to the option whether or not to cede his own 'finite' moral sovereignty to God? His position seems to lead to what we might call a 'sovereignty dilemma', as it is unclear what sort of God would remain if we ourselves supply, or at least vet, its commandments. Either: (i) James permits belief in the existence of a God which countermands his own moral sovereignty and issues demands which might conflict with his own principles, perhaps forcing him to change his principles as those of a mere 'finite demander'; or (ii) the only prescriptions James would allow the God to impose would be ones compatible with his own principles. But where does that leave those 'eternal' and 'overlapping ... things in the universe that throw the last stone'; that 'Thou' who is the 'more perfect and more eternal aspect of the universe'?

James could perhaps respond that surely all reflective believers think there is work to be done in deciding right from wrong according to their religious and/or moral principles. But the issue is precisely what status these other principles have. He seems to allow two incompatible positions: (i) the believer's moral sovereignty, so any other principles she has, could end up subordinate to her God; and (ii) the believer's moral sovereignty stays paramount, so one or more of her other moral principles can take priority over imperatives she believes are from her God.

In the here and now (ii) seems to hold sway: 'There is but one unconditional commandment, which is that we should seek incessantly, with fear and trembling, so to vote and to act as to bring about the very largest total universe of good which we can see' (1891, 2000:259). But projecting forward, James gets closer to (i):

In the interests of our own ideal of systematically unified moral truth ... we ... must postulate a divine thinker, and pray for the victory of the religious cause.

Were this divine thinker to exist,

his way of subordinating [his 'all-enveloping'] demands to one another would be the finally valid casuistic scale; his claims would be the most appealing;

his ideal universe would be the most inclusive realizable whole (1891, 2000:262).

James then tacks back towards (ii):

...exactly what the thought of the infinite thinker may be is hidden from us even were we sure of his existence; so that our postulation of him after all serves only to let loose in us the strenuous mood (1891, 2000:262)

– a mood which, among other things,

makes us quite indifferent to present ill, if only the greater ideal be attained (1891, 2000:260).

In his own case, agnosticism about what the God's imperatives would be protects him from dogmatic and potentially repressive consequences. But his W2B doctrine in isolation is intended as a general 'defence of our right to adopt a believing attitude in religious matters' (1891, 2000:198), without stipulating that others must qualify their religious beliefs in equivalent ways.

Perhaps reflective believers do see a duty to decide right from wrong according to their religious and/or moral principles. But not all believers are reflective, and 'reflective' can mean different things to different people. Whether reflective or not, believers exist who do not share James's agnosticism about the content of divine imperatives – and might think such agnosticism misses the point of religious belief. The concern remains that religious apologetics can lead a believer to stifle her own conscience. JRH evades this by subscribing to a Cowperian divine mystery, but W2B does not mandate it.

If the 'God is sovereign' alternative (i) stands, then James would permit beliefs like the expanded (4.11.1), trumping his tolerance imperative with potentially intolerable results. We should also remember the moral meta-principle of 'universality'. The Golden Rule and Kant's Categorical Imperative for example both entail that whatever is right or wrong, permitted or forbidden, applies to all,

even if the exact domain of 'all' is undefined. For Singer, as for Hume, universality is implicit in moral discourse:

In a dispute between members of a cohesive group of reasoning beings, the demand for a reason is a demand for a justification that can be accepted by the group as a whole. Thus the reason offered must be disinterested, at least to the extent of being equally acceptable to all. As David Hume put it, a person offering a moral justification must "depart from his private and particular situation and must choose a point of view common to him with others; he must move some universal principle of the human frame and touch a string to which all mankind have an accord and symphony." (Hume, 1777, 1902, 1970: Section IX, Part 1, p272; Singer, 2011:93)

In our present context that universality implies that, if W2B allows James to believe in his sort of God, it should allow others to believe in theirs. This does not mean that if it is morally permissible to believe in a certain type of God then it is morally permissible to believe in other forms of God. But the rationale James uses to justify religious belief does not constrain the content of the belief. W2B applies not just to JRH but to any hypothesis forming part of a 'genuine option that cannot by its nature be decided on intellectual grounds' – which JRH would be an example of. His 'tenderness and tolerance' imperative is independent of W2B. If he bases his objection to a principle like CP on the claim that in his specific case, his own religious belief provides moral benefits which outweigh the potential moral benefit of conforming to CP, then although this might count against CP it would not count against EP4 because of EP4's 'conflicting moral imperative' condition. But he seems to consider his generic W2B argument on its own a sound objection to CP. So we are not saying that if we assume it is permissible, *for whatever reasons*, for James to believe in his kind of God then we must assume it is permissible for anyone else to believe in their kind of God. We are arguing that if it is legitimate for James to use his generic W2B argument to justify his belief in his kind of God, then it should be legitimate for anyone else to use that same generic W2B argument to justify their belief in their kind of God. (Conversely: it may or may not be morally permissible to believe JRH, but if it is, the reason is not W2B.)

We would therefore advocate an explicit moral principle like EP4 to guard against over-belief with negative consequences, even if sometimes the principle, like perhaps all moral principles, will lead to moral dilemmas which cannot be resolved by reason alone.

This however opens the possibility of another line of attack: could other moral principles do the same work as EP4; and if so why favour EP4? As this chapter primarily pits EP4 against a contrasting Jamesian position, we shall revisit James's 'sovereignty dilemma' and see if Jamesian alternatives to EP4 can resolve it.

Sovereignty dilemma revisited

Our interpretation of James has generated a dilemma. Either (i) he is defending an individual's right to hold religious beliefs which conflict with his own tolerance ethic.¹⁰⁴ Or (ii) some combinations of religious belief components will fall outside the freedom he would grant. These combinations may not be universal across all religious belief, but they are common and familiar. Of special concern is when belief in a religious dimension to reality ('God' for convenience) incorporates belief in moral imperatives derived from that God, including the imperative to cede one's moral sovereignty to the God one believes exists.

One possible response to (i) is that our interpretation is unfair, and James's position is rather that we are inevitably swimming in a sea of moral and existential choices so the legitimacy of belief only makes sense at the level of (moral) value. But then he could have made his case much quicker, without live hypotheses and genuine options. If his defence of our 'right to adopt a believing attitude in religious matters' (1896, 2000:198) was just that, in respect of a specific religious belief, it would be legitimate if it led to a better moral outcome than not believing, then he only need argue there are occasions when it is right to disobey CP, just as there are occasions when it is right to lie. EP4 recognises this. Also it is not obvious that

¹⁰⁴ See p185ff.

when it is right to disobey CP would necessarily be when the individual is faced with a Jamesian genuine option.¹⁰⁵

Arguably, James's treatment of live and dead hypotheses makes horn (i) even more problematic. He lets an individual believe a live hypothesis as an act of will if it forms part of an undecidable genuine option. But a hypothesis can be live as an indirect result of an act of will. So if there is nothing (including no conflicting evidence) to make the hypothesis dead, one or more actions of her 'willing nature' can indirectly make it live. She may not be able to make it live by direct volition, but could, by deliberately suppressing doubt and controlling her social interactions and intellectual exposure, nurture it into life. She could then legitimately believe it, as long as it forms part of a genuine option.

So if we take horn (i), James would allow the individual to end up legitimately holding a set of descriptive and prescriptive beliefs with intolerant and repressive consequences, for example to persecute blasphemers, if that is what she believes her God commands.

A critic might object that if the set of live hypotheses generates a dilemma – and/or supposedly legitimate beliefs with morally objectionable consequences – that would count as 'intellectual grounds' for non-belief. But believers who insist on holding intolerant and oppressive versions of hypothesis (4.11.1) are not obviously guilty of a logical contradiction or factual mistake. For example I believe a religious hypothesis in the form of (4.11.1) can have negative moral consequences which would be reason enough to reject it. This is because I see a supreme metaphysical belief fused with a supreme prescriptive belief as dangerous. But this is a moral view, not a theoretical position. A theist trusting her God's

¹⁰⁵ It could be claimed that on James's overall position it is permissible to believe a genuine option provided that belief itself will not have clear negative moral consequences, whereas on EP4 it is only permissible to over-believe if there is a moral benefit. On the face of it this seems a distinction, with James's position being more permissive. But, as mentioned previously (pp153–156), we should not make too much of the apparent difference between EP4's requirement of a net moral benefit versus James's requirement of no net moral harm. In exceptional circumstances EP4's 'conflicting and overriding moral imperative' condition could be achieved purely by preserving an individual's right to liberty where no harm is likely to be incurred by a third party. We need to remember the rationale for EP4, as that for CP, in avoiding the insidious and corrosive effects of over-belief in a social context, which James by comparison seems to downplay. From the perspective of EP4 and CP the circumstances when over-belief incurred no third-party risk would generally be seen as exceptional.

goodness may not share my concern. Had James also seen this fusion as dangerous he would have expressed W2B and/or JRH differently.

We could perhaps loop round again and suggest this may reinforce James's point since what lies at the bottom is a moral judgment, not a factual belief. But then James and EP4 are not in conflict, as EP4 excludes prescriptive beliefs. And besides, is James really saying it rests on a moral decision? His position seems more permissive than saying we are morally constrained as to acquiring and holding descriptive beliefs: that for example we should generally meet an evidential requirement, but there are circumstances when the possibility of net moral benefit overrides this imperative.

James could argue he has been clear which horn he would take:

I do not think that any one can accuse me of preaching reckless faith. I have preached the right of the individual to indulge his personal faith at his personal risk. I have discussed the kinds of risk; I have contended that none of us escape all of them; and I have only pleaded that it is better to face them open-eyed than to act as if we did not know them to be there. (1897, 2000:195)

But his approach to religious hypotheses is revealing:

If religious hypotheses about the universe be in order at all, then the active faiths of individuals in them ... are the experimental tests by which they are verified, and the only means by which their truth or falsehood can be wrought out. The truest scientific hypothesis is that which ... 'works' best; and it can be no otherwise with religious hypotheses. Religious history proves that one hypothesis after another has worked ill, has crumbled at contact with a widening knowledge of the world, and has lapsed from the minds of men. Some articles of faith, however, have maintained themselves through every vicissitude ...: it is for the 'science of religions' to tell us just which hypotheses these are. Meanwhile the freest competition of the various faiths with one another, and their openest application to life by their several

champions, are the most favorable conditions under which the survival of the fittest can proceed. ... [I]t is only when they forget that they are hypotheses and put on rationalistic and authoritative pretensions, that our faiths do harm. The most interesting and valuable things about a man are his ideals and over-beliefs. (1897, 2000:196-7)

So even if a hypothesis is believed, it stays a hypothesis. A religious hypothesis, like a scientific hypothesis, need only 'work'. We should therefore be reassured that James's own religious hypotheses would never 'forget that they are hypotheses'. But what of hypotheses which do forget, and which include imperatives which (4.11.1) could include? Is it morally safe or sound to leave these and their 'several champions' to fight it out to see which get 'verified' and which 'crumble' or 'lapse'? Ideals and over-beliefs can also be our most dangerous things, inflaming a survival of the fittest into a Thirty Years War.

The concern is that, to be palatable, James's doxastic permissiveness – considering W2B in isolation – must be constrained by other moral principles (for toleration, say, and against 'rationalistic and authoritative pretensions'), which would then stop religious belief submitting to a 'more perfect and more eternal... Thou' (1896, 2000:216), an 'infinite demander ... with all-enveloping demands' (1891, 2000:262) – which for James seems the whole point of religious belief.

A critic might object that the moral optimism which JRH projects does not imply there cannot be moral assessment of religious claims. But James has two separate theses: W2B, leading to a conditional, but generic, doxastic permissiveness; and his specific JRH, which includes his moral optimism. JRH is an example of what he thinks an individual has the right to believe, but this does not limit what W2B would allow others to believe. JRH speaks of an 'infinite demander ... with all-enveloping demands', but also incorporates agnosticism about what that infinite demander's demands are as they would be 'hidden from us'. But there is no necessity about that agnosticism: W2B would let another individual subscribe to another 'infinite demander' and not only think she knows its demands, but also make them sovereign over her own and everyone else's moral values. We are not claiming anything leads to a situation where there cannot be

moral assessment of religious claims, but that W2B can let an individual believe there cannot be (individual, human, 'finite') moral assessment of religious claims. W2B does not seem to be a thesis which can be left unqualified.¹⁰⁶

And although James's picture of natural selection among 'hypotheses' and 'faiths' seems not to reflect reality (where the 'several champions' can forget their hypotheses are only hypotheses) we must assume he intended his comments to address the real world. So if one is balancing risk that would be another reason to back a principle like EP4 against Jamesian permissiveness – which is not just about what hypotheses to consider, but about what and how to believe.

He is, after all, asserting our right to '**decide** an option between propositions' (1896, 2000:205), our 'right to adopt a **believing** attitude in religious matters' (1896, 2000:198), our 'right to **believe** at our own risk any hypothesis that is live enough to tempt our will' (1896, 2000:217). He hardly needs the length of his essay to argue our right to entertain a hypothesis, something merely '**proposed** to our belief' (1896, 2000:199). [Emphases added]. Neither Clifford nor EP4 would deny this right. Perhaps therefore we should see his talk of hypotheses versus rationalistic faiths as distinguishing not between entertaining something and believing it but between fallible belief and dogmatism.¹⁰⁷

We will now set aside worries about apparent inconsistencies within James's overall position, and evaluate two similar but related alternative moral principles,¹⁰⁸ both appropriately 'Jamesian' and 'liberal', to see if either or both could be as effective as EP4 in ensuring we avoid legitimising harmful prescriptive beliefs.¹⁰⁹

¹⁰⁶ In response to a potential 'tu quoque' we should point out that W2B is formulated to identify a distinct category of beliefs representing a principled exception to an evidence principle like CP, derived from James's conception of a 'genuine option'. Our objective has been to see if there is anything in W2B to threaten EP4, which we are offering as a candidate moral principle based on a modified wording of CP, and with a more restricted application than CP. The wording of EP4 already reflects some potential conflicts, but in practical applications EP4, like every other moral principle, will always encounter other conflicts because EP4, like most if not all moral principles, is not an absolute, in that it does not automatically override other moral principles. Our claim is simply that W2B is less successful than James seems to have thought in identifying a distinct category of beliefs representing a principled exception to a principle like CP.

¹⁰⁷ See also pp161–163.

¹⁰⁸ See p190.

¹⁰⁹ A critic could perhaps complain this is long overdue, and that there has all along been a reasonable and reasonably clear way to interpret James. This is that we are permitted to believe

Alternatives to an evidence principle

We are taking it as given that intuitively immoral views can sometimes be supported by religious beliefs. One way to invalidate such views would be a principle like EP4. But we could instead try setting other moral principles against them.

For example a Jamesian ‘tolerance principle’ (1902, 1960:489-90):¹¹⁰

WJT One must always be tolerant of others’ beliefs which seem to contradict views of one’s own held on insufficient evidence.

WJT is clearly socially contextualised. We could also though propose an equally Jamesian, non-socially contextualised, ‘fallibility principle’:¹¹¹

WJF It is always wrong to hold beliefs for which there is insufficient evidence dogmatically.

We will now assess the effectiveness of these principles against someone who believes her moral beliefs represent the will of the God she believes exists and is the source of morality.

A potential issue with both is that once you have made the leap into religious belief and then into believing your God supreme over your own moral conscience, you could feel justified in believing your God decides what counts as tolerance or dogmatism. This will not of course always apply: the concern though is whether

either hypothesis in a genuine option, providing the resultant belief does not have clear negative moral consequences and is held non-dogmatically. It seems however at least equally reasonable to pursue, as we have done, the implications of a religious hypothesis which is part of that genuine option but which describes a divine source of moral law (‘infinite demander with all-enveloping demands’) with a different evaluation of what the positive and negative moral consequences of the resultant belief might be and which insists that that same resultant belief is not open to question.

¹¹⁰ See p181.

¹¹¹ [T]he intellect, even with truth directly in its grasp, may have no infallible signal for knowing whether it be truth or no (1896, 2000:208).

WJT or WJF could ever let through a belief the content of which could then neutralise either or both principles.

Take WJT. If you believe your God is the supreme God, must you be tolerant of someone else saying their God is the supreme God? A lot will hinge on the exact meaning of 'tolerance'. If you believe H could you be 'tolerant' of someone who believes not-H? Perhaps by not doing anything which would count as denying the other's right to believe not-H? For example you should not vote to make the other person's belief in not-H illegal, or persecute her for expressing her belief in not-H. But if you really believe your God exists and is the unique creator of the universe you cannot 'tolerate' another's view that her God is the unique creator of the same universe, if this means holding two contradictory beliefs at the same time.

And it might be hard to hold a principle of fallibility like WJF at the same time as a belief in a God claiming to be the infallible supreme being and creator of the universe. Not only must the God's imperatives be subordinate to WJT and WJF, but the God cannot impose its own criteria of tolerance and dogmatism.

For a more secular concern with WJF, there are times when the last thing you would want is a principle of fallibility. As you hurl yourself across a chasm would you want to remember you could be wrong or would you rather believe as dogmatically as possible that you will reach the other side? Indeed it might be psychologically impossible to reap the benefits of over-belief while staying nondogmatic. Compare this with EP4 where the justification for your over-belief is your conflicting and overriding moral imperative to stay alive – for your loved ones if not for yourself.

James may not accept this implication of WJF as he uses the chasm story against Clifford.¹¹² Having convinced himself that his leap will succeed James might cling to his conviction 'as if it never could be reinterpretable or corrigible' (1896, 2000:207), with the cake he has just eaten still safe in his hand.¹¹³ His defence could be that it is a 'genuine option' and he was believing dogmatically at his 'own

¹¹² See p112.

¹¹³ See p157.

risk'. That in turn suggests an 'own risk' condition should be added to WJF if not WJT as well.¹¹⁴

We come now to perhaps the most serious apparent weakness of WJT and WJF compared to a more explicit evidence principle. In their basic form they both largely ignore the social implications of Clifford's justification for CP. Ostensibly we might appear to be just over-believing but inwardly exercising tolerance and fallibility. We therefore risk encouraging others to over-believe.

You could of course hold both EP4 and WJT or WJF simultaneously, or indeed all three. They do not contradict each other. We are not arguing WJT and WJF are wrong, just that neither one nor both together seem powerful enough yet to replace EP4. Could we though make the case that having EP4 does not remove the rationale for WJT or WJF, as even obeying EP4 you could still end up holding an over-belief, arising from a conflicting and overriding moral imperative? You could for example be convinced, with no counter-evidence to suggest otherwise, that because of something in your make-up you need to believe a God exists in order to be the best person you can be. EP4 would then allow (indeed require) you to believe in your particular God.

Perhaps WJF might then stop you holding your belief dogmatically, by encouraging you to be open to alternative ways of thinking which might lead you to re-evaluate your belief and/or your reasons for believing. But EP4 on its own should do this, as it only exonerates over-belief while the conflicting and overriding moral imperative applies.

As far as WJT is concerned, there is little wrong with a general principle of tolerance, but again no particular reason to think EP4 incomplete without it. A justification someone holding EP4 might have for holding an over-belief in God is to be the best person she could be. This would typically entail conforming to other moral principles: tolerance, yes; but also love, justice and so on.

¹¹⁴ We will pick up on this later: p198.

Would WJT and/or WJF be as effective as EP4 against the potential negative consequences of over-belief in general and of religious over-belief in particular? This seems unlikely: if, that is, you broadly accept the Cliffordian social rationale for CP. If you do accept this, then you would in any case need to add James's 'at your own risk' into the mix.

We are approaching a position where the attempt to distinguish the practical consequences of adopting EP4 from the practical consequences of a Jamesian alternative gets difficult.¹¹⁵ Clifford emphasises the negative social consequences of over-belief while James downplays them. But (we must ask although it may be unanswerable) what if both were privy to the same sociological evidence and saw that evidence in the same way?

We are characterising the Jamesian position as essentially liberal: stressing individual liberty but constrained by social impact. The contrasting Cliffordian position could be described as socially protectionist. In drafting EP4 we have modified this, as we have other aspects of CP, but without sacrificing its core content.

The difference between James and Clifford/EP4 can be illustrated by what seem their respective motivations. James appears driven by a desire to preserve religious belief as a legitimate endeavour, Clifford by a desire to optimise epistemic behaviour and minimise the negative social consequences of endemic credulity. EP4 shares Clifford's priority, but offset with a desire to maintain fairness. EP4 does not share James's aim of preserving the legitimacy of religious belief, but it aims to be morally impregnable. The EP4 position would see no intrinsic value in religious belief itself has, but a person's right to hold a religious belief if this is truly at her own risk could have an intrinsic value which any EP4 'verdict' should reflect.

It is possible therefore that a Jamesian and a proponent of EP4, faced with the same evidence concerning the consequences of doxastic behaviour and with the

¹¹⁵ Alluded to earlier p181.

same evaluation of that evidence, could in practice make the same judgments: the Jamesian from a desire to legitimise religious belief as long as it avoids negative consequences; the EP4-proponent from an imperative to promote good doxastic behaviour without restricting individual freedom unnecessarily, in areas where no harm will result or it is very unlikely that harm will result. It would not obviously defeat the purpose of EP4 if the doxastic liberty the Jamesian says you have a right to when any risk is truly only to yourself would count for EP4, all by itself, as a conflicting and overriding moral imperative.

It is also worth noting that WJT and WJF are themselves worded as evidence principles of a sort. Like EP4 they both specify what should be done in respect of beliefs held on insufficient evidence. All three permit beliefs held on sufficient evidence without further condition. So both WJT and WJF, like EP4, presuppose a workable criterion of 'sufficient evidence' – which will be the subject of **Chapter 5**.

The Jamesian cluster of WJT + WJF differs however from EP4 in their treatment of over-beliefs:

WJT: Over-beliefs are legitimate as long as one is tolerant of others' over-beliefs which contradict them.

WJF: Over-beliefs are legitimate as long as they are held on the assumption of fallibility, so not dogmatically.

EP4 by comparison considers over-beliefs legitimate as long as they are necessitated by a conflicting and overriding moral imperative.

Neither WJT nor WJF as currently worded addresses the potential social consequences of over-belief which EP4, following CP, targets in its focus on doxastic behaviour. To plug this gap we added the Jamesian condition of 'at the believer's own risk'. So we can now summarise an overall Jamesian position of conditional freedom:

Over-beliefs are legitimate as long as the conditions of tolerance (WJT), fallibility (WJF) and 'own risk' are met.

EP4 is currently expressed in terms of mitigated constraint, but could be recast in terms of freedom:

Over-beliefs are legitimate as long as there is a net moral benefit in allowing yourself this freedom (and/or a net moral cost in not allowing yourself this freedom).

EP4 effectively wraps up the three 'Jamesian' conditions in its single blanket moral condition.

Now to reintroduce the Cliffordian 'shared epistemic asset' in the rationale for EP4 and its focus on doxastic behaviour. Ostensibly the Jamesian stance does not accommodate this. To ensure equivalence we should therefore stipulate that the 'own risk' condition incorporates this: to reflect the evident possibility that poor doxastic behaviour could have a negative impact on our shared epistemic asset.

We mentioned earlier¹¹⁶ that, in a scenario where EP4's conflicting and overriding moral imperative permitted an over-belief, it would be because it was obligatory. This contrasts with James's 'genuine option', where the over-belief would be permitted but not obligatory. Our comparison between the Jamesian position and EP4 provides another explanation of this difference. EP4 presupposes a default duty to protect our shared epistemic asset. This means EP4 would exclude the freedom to choose between over-belief and disbelief, since in any context where disbelief would be permitted the over-belief alternative would not be open to us as, other things being equal, over-belief endangers our shared epistemic asset.

With this exception we seem to have reached a position of almost equivalence. But we can still contrast the Jamesian stance with EP4 in that the Jamesian ultimately prioritises individual freedom whereas EP4 ultimately prioritises our

¹¹⁶ p184.

shared social epistemic asset. In some ways this mirrors a political choice between liberty and social security. In a political context individual freedom needs protection by the same social structures as can curtail it; and our shared epistemic asset plays a key role in holding those social structures together.¹¹⁷ And indeed if the choice between EP4 and a Jamesian position boils down to a moral dilemma then that in itself supports the case that EP4 qualifies as a bona fide moral principle. Otherwise EP4 would not feature in a moral dilemma of this kind.

4.5 Conclusions

We will first summarise where the previous section has brought us, and then set that in the context of **Chapter 4** as a whole.

Summary of previous section

We introduced James's religious hypothesis (JRH)¹¹⁸ as an example of what his will-to-believe thesis (W2B) would permit, not as the only thing it would permit. JRH, like our (4.1.1–3),¹¹⁹ has both descriptive and prescriptive components.

We then evaluated his defence of religious belief like JRH as arising from a genuine option¹²⁰ and therefore legitimate. We doubted to what extent the purely descriptive content of a hypothesis like JRH could qualify as belonging to a forced option, and thought it more plausible that its prescriptive content was what made a religious belief option momentous and therefore forced. In religious belief though the descriptive content can make the prescriptive content binding. In some scenarios this could lead to something approaching circularity,¹²¹ which we can avoid by applying a principle like EP4.

JRH shares another familiar feature of religious belief, in that it characterises the posited God as a person¹²² rather than an impersonal entity. This is significant

¹¹⁷ See however p294ff.

¹¹⁸ pp166–167.

¹¹⁹ p131.

¹²⁰ pp167–179.

¹²¹ p177.

¹²² pp179–181.

because interpersonal contexts can introduce further doxastic duties favouring belief over disbelief.

James insists however that religious belief is essentially private so at the believer's 'own risk'.¹²³ This confronts Clifford's rationale for CP based on the potential harm encouraged habits of credulity can cause. James's counter-proposal is the unique moral benefit religious belief can provide. We concluded that a debate between James and EP4 ultimately boils down to how to resolve a moral dilemma, which suggests the distinction between James and EP4 is more of stress than of principle.

The next move however indicates a substantive contrast. This is what W2B could imply for the believer's 'moral sovereignty':¹²⁴ what if the believer believed her God insisted its imperatives superseded any other principles she might have; particularly when its imperatives were repressive or otherwise socially harmful? We concluded that James's overall position appears inconsistent, in that W2B could allow an individual to hold religious beliefs which might trump one or more of James's other stated principles, for example advocating tolerance.

This however assumes all relevant parts of James's stance on religious belief are intended to be consistent with each other and to be taken at face value. As the overall thrust of **Chapter 4** is to test EP4 against James's alternative position, we tried a more generous interpretation to see if his writings could provide alternative moral principles which might do EP4's job, at least as well if not better.

After revisiting the 'sovereignty dilemma'¹²⁵ we proposed two distinct but related principles, promoting tolerance and fallibility respectively.¹²⁶ We concluded that an essentially Jamesian composite position could be constructed which operates almost as successfully as EP4, particularly if we include the Jamesian condition of 'at the believer's own risk'. Although the Jamesian composite and EP4 will still have features revealing their different origins, the practical consequences could be

¹²³ pp181–185.

¹²⁴ pp185–190.

¹²⁵ pp190–195.

¹²⁶ pp195–201.

virtually identical – particularly if we assume the same evidence of social consequences and the same evaluation of that evidence.

Although the Jamesian position is fundamentally liberal while EP4 (following Clifford) is more social-protectionist, in a practical (and political) sense the originating values may be inseparable, as social structures need protection to guarantee individual liberty.

Broader context and chapter conclusions

Our overall conclusion is that James would probably reject EP4, but, assuming we can address its as yet undefined notion of ‘sufficient evidence’,¹²⁷ EP4 should survive whatever he might throw at it. If his position on religious belief is incompatible with EP4, then the choice between them is a moral one. Nothing in his position invalidates EP4 as an independent moral principle, as EP4 is not an absolute over other principles. Although W2B asserts a domain of permission which EP4 denies, in the context of religious belief the only circumstances where his ‘right to believe’ would override the moral wrong of over-belief would be where other moral considerations make religious over-belief obligatory and therefore permitted.

It might be objected that surely there are situations where both evidence and morality leave it open what to do. There are after all many ways to resolve a moral dilemma. But we are only talking here about deciding what and how to believe and, specifically, deciding which doxastic behaviours to nurture and which to avoid and discourage. We are not discussing moral dilemmas in general but circumstances in which adopting a principle like EP4 would systematically lead to a net moral loss. We can resolve moral dilemmas in many ways, but this chapter is about moral dilemmas specific to the acquisition and retention of descriptive religious beliefs.

In saying any incompatibility between EP4 and Jamesian permissiveness comes down to a moral choice, we would claim EP4 as morally superior: equivalent to a

¹²⁷ See p163.

Singer or a Bentham judging an imperative to consider the potential net happiness of all sentient beings morally superior to an imperative considering only the potential net happiness of all humans. EP4 presupposes a moral context ('If anything is morally wrong, then...') where moral decisions and standpoints can be evaluated and compared against each other. We are claiming EP4 is morally superior to Jamesian permissiveness for all the reasons this chapter provides. We cannot however *prove* it is morally superior, but nor could we prove the Singer/Bentham verdict was morally superior to a purely anthropocentric alternative.

We could give reasons in favour of Singer/Bentham, for example that sentient beings other than humans feel pain, but this may not convince anyone who did not care whether non-human sentient beings felt pain. The move from (generic) 'allow religious over-belief' to (more generic) 'do not allow over-belief' involves an increase of dimension: from seeing the issue as only impacting the individual to one addressing all potential social consequences. The move from (generic) 'consider potential net happiness of all humans' to (more generic) 'consider potential net happiness of all sentient beings' involves an analogous expansion of dimension. But one could remain unmoved by this 'expanding circle'.¹²⁸

¹²⁸ See Singer (2011). Singer got his title from Lecky (1880: 100-101).

Chapter 5 | Sufficient and insufficient evidence

5.1 Introduction

We began with Clifford's principle:

CP *[I]t is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence (1877, 1879:186).*

Clifford's 'anything' though includes categorical prescriptive beliefs (for example moral beliefs), which we found compelling reasons to exclude.¹

So we modified CP and tested it against apparent counter-examples, from involuntary and innocuous over-beliefs to beliefs about the future and the past and beliefs about personal relationships.² The final result was a more qualified and explicitly moral principle:³

EP4 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything [knowingly or irresponsibly] on insufficient evidence {in the absence of any conflicting and overriding moral imperative} [except when the unjustified believing is outside the believer's voluntary control].

We then discussed James's 'defence of our right to adopt a believing attitude in religious matters' (1896, 2000:198).⁴ Clifford did not provide a satisfactory account of 'insufficient' (or of 'sufficient') evidence,⁵ but we concluded that, as long as we can fill this gap, EP4 should withstand Jamesian attacks.⁶ James would probably consider this a 'fool's errand' (Hall, 2011:82):

¹ See pp43–73. We do not however need to exclude hypothetical or 'prudential' prescriptive beliefs: see pp44–46.

² See p74ff.

³ p92.

⁴ See p129ff.

⁵ See p163.

⁶ See p201ff.

No concrete test of what is really true has ever been agreed upon. ... The much lauded objective evidence is never triumphantly there; it is a mere aspiration or Grensbegriff [sic]⁷ marking the infinitely remote ideal of our thinking life. To claim that certain truths now possess it, is simply to say that when you think them true and they are true, then their evidence is objective, otherwise it is not. But practically one's conviction that the evidence one goes by is of the real objective brand, is only one more subjective opinion added to the lot. (James, 1896, 2000:207-208)

James appears to conflate 'objective' with 'conclusive' evidence,⁸ so his complaint that Clifford regards 'objective evidence' as something guaranteeing 'bottom-certitude' seems unfair. 'Objective' does not even appear in 'The Ethics of Belief': 'common property'⁹ is something different.¹⁰ But this chapter's task is to find a suitable criterion of evidential sufficiency.

Many explanations have been proposed of how evidence supports belief. Without space to review them all we will see what a leading candidate – Bayesianism – can offer.

Bayesianism should be of particular interest because at least some versions offer prima facie support for a Jamesian kind of 'voluntarism' typically contrasted with 'evidentialist' principles like EP4.¹¹ Also as an explicitly moral principle EP4 presupposes an evaluative context rather different from ones Bayesianism usually inhabits, with the latter's focus on rationality, which for Bayesians requires conformance to the probability calculus. On the other hand Bayesianism has an evident quantitative rigour and intuitive plausibility particularly relating to changes in degree of belief. Our aim is not to critique Bayesianism, but see if it contributes

⁷ Limit or ideal notion.

⁸ See p161.

⁹ A chemist's 'result becomes common property, a right object of belief, which is a social affair and matter of public business.' (1877, 1879:197)

¹⁰ See also **5.6 Indefinite intersubjectivity** p253ff.

¹¹ See Van Fraassen (1989:Chap. 7) for the development of how the Bayesian framework can be part of a sophisticated voluntarist epistemology. Van Fraassen does however describe himself as 'a probabilist, though not a Bayesian' in that he 'do[es] not accept the Bayesian recipes for opinion change as rationally compelling' (1989:175).

to an account of evidential sufficiency for EP4, without expecting it to provide all the answers.

We first dig deeper into what kind of relationship a criterion of sufficient evidence could have with a principle like EP4, before expanding on EP4's evaluative context in line with Clifford's approach to belief. We then introduce Bayes' Rule, and the way it quantifies how increasing evidence for (or against) a hypothesis supports an increased (or decreased) 'personal probability' – or 'degree of belief' or 'credence' – for that hypothesis.¹² Applying Bayes' Rule to a test case¹³ introduces suggestions from different variants of Bayesianism as to the source of our initial credences ('prior probabilities'), and what may constrain them. Bayesianism eventually helps us understand 'insufficient' evidence but not yet 'sufficient' evidence.¹⁴

We next ask if any normativity involved in a Bayesian approach comes just from conformance to the probability calculus or involves implicit acceptance of normative decision theory.¹⁵ We conclude that in the context of EP4 we can justifiably assume the latter.

We then address the transition from high credence to full justified belief,¹⁶ and consider the idea of a quantitative threshold above which credence may count as belief. However issues arising from 'lottery cases'¹⁷ cast doubt on how Bayesian 'imaginary gambles' operate around this threshold. Qualitative differences in a believer's context¹⁸ and in what a believer may have at stake¹⁹ also appear relevant to justified belief versus mere high credence. Both represent ways of factoring 'practical interests'²⁰ into apparently justified belief which are at least compatible with Bayesianism, particularly in our assumed context of normative decision theory.

¹² p213ff.

¹³ p216ff.

¹⁴ pp224–226.

¹⁵ p226ff.

¹⁶ p230.

¹⁷ p233.

¹⁸ p244ff.

¹⁹ p240.

²⁰ p246ff.

A potential issue for EP4 though is whether considering individual practical interests is compatible with an evidence principle expressed in moral terms.²¹ Our suggested resolution is to extend the moral envelope to include those practical interests themselves: if we enlist practical interests to justify a belief then for EP4 as a moral principle they should be the practical interests we (morally) ought to have, not just ones benefitting us as individuals.²² Bayesian 'imaginary gambles' can have stakes of social rather than purely personal utility.

For EP4 as for Clifford, the practical interests we ought to have will include the preservation and enhancement of our shared social asset of knowledge and true belief. This helps neutralise the objection that the notion of practical interests could jar with the requirement of 'objective' evidence we would expect from a principle like EP4.

The next step is to see if anything in Craig's 'state of nature' epistemology can help us understand 'sufficient evidence', aligned with the 'practical interest we ought to have'. Craig originates an 'objectivised' concept of knowledge in the idea of a 'good informant ... whatever the particular circumstances of the inquirer' (1990, 1999:91). We argue that something very like this holds promise for the criterion of sufficient evidence which EP4 needs.

5.2 How a criterion of sufficient evidence relates to EP4

A moral principle like EP4 can be analysed into more than one component. The specifically moral 'thou shalt/shalt not' concerns the believing itself. On the face of it the criterion of evidential sufficiency does not itself need to be a moral component. To this extent EP4 resembles other moral principles.

But the moral claim in 'it is wrong to deceive' is about the wrongness of deceiving, not about what counts as deception, which the moral claim presupposes. Similarly 'it is wrong to lie' says it is wrong to say something false, but does not specify

²¹ p249ff.

²² See also p295.

what makes an utterance false. EP4 on the other hand does not just say it is wrong to hold or acquire an unjustified belief. It says that, exonerating conditions aside, it is wrong to believe except on sufficient evidence.

Evidence itself can count as belief(s), if not knowledge. A condition of sufficient evidence includes that of consistency (coherence) between beliefs, and therefore has an element of self-reference. We will see how the Bayesian approach articulates this explicitly.²³

Contrasting 'sufficient' with 'insufficient' implies a judgment, and therefore a (potentially moral) requirement to exercise judgment. If the condition of sufficient evidence is at least partly about coherence between beliefs, then part of that sufficiency would be sufficient coherence. An agent could believe hypothesis H on insufficient evidence by failing to realise H was inconsistent with his other (evidential) beliefs E_{1-n} . But he could also fail to acquire evidential beliefs E_{1-n} , which might instead support not-H.²⁴ Believing on insufficient evidence could involve either failing. To ensure he obeys EP4 he must exercise judgment about where to draw the line between 'sufficient' and 'insufficient' in both respects.

Again we can compare other moral principles. To ensure you are not causing unnecessary suffering takes judgment about where suffering starts and stops. To avoid lying takes judgment about the difference between truth and falsehood. These judgments also incorporate judgment about how to apply judgment. If you think what you are doing might cause pain to another then you should decide as quickly as possible whether to stop or continue.

Moral principles like EP4 seem therefore to be analysable into three components:

- (i) A moral 'thou shalt/shalt not'.
- (ii) A non-moral 'domain' or subject matter (utterances in the case of lying; something like 'sentience' in the case of suffering; other (evidential)

²³ See p219ff and p222ff.

²⁴ Expanded later in Bayesian context p227ff.

beliefs in the case of EP4).

- (iii) A moral requirement to exercise judgment in the context of domain (ii).

An evidence principle could be thought to differ from other moral principles by concerning just the individual believer, and governing behaviour not directed towards others. Clifford however would disagree. We will review his reasons, which also apply to EP4.

5.3 Social context of belief: 2

Clifford stresses the social, and therefore moral, context of believing.²⁵ One person's beliefs and ways of believing only make sense in the context of other people's beliefs and ways of believing. A community's beliefs and ways of believing aggregate into an 'heirloom' (1877, 1879:182) which successive generations inherit, enhance, critique, modify and bequeath, helping 'create the world in which posterity will live'. James expresses similar sentiments:

Our fundamental ways of thinking about things are discoveries of exceedingly remote ancestors, which have been able to preserve themselves throughout the experience of all subsequent time. (1907, 2000:Lecture 5)

Clifford however sees socially normative implications. What and how we believe 'is ours not for ourselves, but for humanity' (1877, 1879:183), because it helps 'bind men together' and 'strengthen and direct their common action'. Our duty to this social asset, and to others in respect of our shared custodianship, is 'hard' and shared by all.²⁶ It entails a further shared duty to avoid the 'savagery' which a 'credulous' society is prey to.²⁷ The more we protect our shared epistemic asset

²⁵ See p16ff and p35ff.

²⁶ See p16.

²⁷ See p16.

and discourage credulity, the more we can rely on testimony.²⁸ To believe what a person says, we need

reasonable grounds for trusting his veracity, that he is really trying to speak the truth so far as he knows it; his knowledge, that he has had opportunities of knowing the truth about this matter; and his judgment, that he has made proper use of those opportunities in coming to the conclusion which he affirms. (1877, 1879:189)

A scientist's authority is valid because others 'question it and verify it', keeping 'alive among investigators the love of that which shall stand all possible tests' (1877, 1879:196-8). But it is not only scientists who should question and verify. The 'labours and struggles' of our predecessors have bequeathed an 'atmosphere of beliefs and conceptions', of 'forms and processes of thought', allowing us to think and 'breathe amid the various and complex circumstances' of our lives (1877, 1879:199). It is 'not only possible and right, but our bounden duty' to 'doubt and to test' this inheritance. Again we see a parallel in James, but while James offers pragmatic advice to the individual ('Retain, I pray you, this suspicion about common sense' (1907, 2000:Lecture 5)), Clifford stresses social duty:

[T]he main purpose of the tradition itself is to supply us with the means of asking questions, of testing and inquiring into things; [so] that if we misuse it, and take it as a collection of cut-and-dried statements, to be accepted without further inquiry, we are not only injuring ourselves here, but by refusing to do our part towards the building up of the fabric which shall be inherited by our children, we are tending to cut off ourselves and our race from the human line. (1877, 1879:199)

The 'simple rule' is to use our inheritance as its 'makers' used it: 'to ask further questions, to examine, to investigate' (1877, 1879:203). We therefore pursue good inquiry as well as knowledge. Our shared asset contains not only facts and true or at least justified beliefs, but also 'questions rightly asked', conceptions which

²⁸ See also pp91–95.

‘enable us to ask further questions’, and ‘methods of answering questions’. Its ‘very sacredness’

imposes upon us the duty and the responsibility of testing it, of purifying and enlarging it to the utmost of our power. He who makes use of its results to stifle his own doubts, or to hamper the inquiry of others, is guilty of a sacrilege which centuries shall never be able to blot out. (1877, 1879:205)

CP can apply to an individual belief or act of believing. More generally it summarises our duties in respect of our shared ‘heirloom’ – which combines knowledge, true and/or justified beliefs, ways of questioning, and principles for its own onward critical custodianship.

EP4 is more constrained and less universal than CP but equivalently contextualised, with ‘morally’ and ‘moral’ appearing three times in total. In theory a moral context need not imply a social context. But we are explicitly excluding that possibility. The rationale behind both CP and EP4 is the potential social consequences of evasion. So while EP4 can apply at the level of an atomic belief or act of believing (assuming these are possible), its scope also covers shared duties in respect of our shared epistemic legacy.

We should remember this when discussing Bayesianism. The Bayesian believer is generally characterised as an ideally rational agent, where to be ‘rational’ is either to be merely coherent (in the sense of conforming to the probability calculus²⁹) or to seek to maximise individual personal utility. The social and moral context presupposed by both EP4 and CP would seem *prima facie* more demanding than that of a collective of ideally rational Bayesian agents.

One last preliminary point. CP and EP4 both address belief rather than knowledge, allowing us to avoid debates about when and how justified belief qualifies as knowledge. In some contexts however it will be convenient to drop into

²⁹ See p214.

'knowledge' language. Although not every justified belief qualifies as knowledge,³⁰ in general the sufficient evidence a justified belief needs is also what it needs to qualify as knowledge. 'Knowledge' language can therefore be a useful proxy for 'justified belief' language, and will sometimes be used as such in what follows. But we are not rephrasing or extending EP4 to speak of knowledge rather than belief.

5.4 Bayesian approaches

We will not attempt a derivation of Bayes' Rule³¹ and its founding principles. We will take these as given and see if they help clarify 'sufficient' and/or 'insufficient' evidence for EP4.

Notions of evidential sufficiency or insufficiency are not strictly part of Bayesianism. But Bayes' Rule does offer 'a way to represent rational change in belief, in the light of new evidence' (Hacking, 2001, 2009:171). Key Bayesian tenets are that a 'degree of belief' ('credence') is a subjective or personal probability,³² and as such must conform to axioms of the probability calculus or 'standard mathematics of probability' (Strevens, 2005:2-3).³³ The Bayesian account of rational change in belief rests on a 'rule of conditionalization'³⁴ (Hacking, 2001, 2009:259).

One way to apply the probability calculus to 'personal degrees of belief' is by representing them 'numerically by using imaginary gambles' (2001, 2009:150). A personal degree of belief can be expressed as a number between 0 and 1 representing a personal probability or 'betting rate'.³⁵

An actual or candidate belief (whatever degree of belief it is held at) is a 'hypothesis' or a proposition:

³⁰ Articulated further p264.

³¹ To avoid ambiguity we will refer to "Bayes' Rule" rather than "Bayes' Theorem": see Strevens (2005:4-5).

³² Expanded later p226.

³³ The next subsection will unpack this conformance in the context of an example: p219ff.

³⁴ See p215, and also p223.

³⁵ Although '[b]elief is not the sort of thing that can be measured exactly' (Hacking, 2001, 2009:155-6).

When you think about your personal probability for a hypothesis, you do so relative to your background knowledge, beliefs, prejudices, and so on.
(2001, 2009:172)

A Bayesian sees a degree of belief (credence) as something you may have without necessarily being aware of its strength relative to your other credences. To make that relative strength evident, a Bayesian may appeal to an imaginary gamble where you are asked for a fair betting rate such that you would be equally happy to take either side of the bet. The following articulates your personal probability p that hypothesis H is true:

If p is, in your opinion, the fair rate for betting on [H], then you should be indifferent between:

A bet on [H] at rate p : you win $(1 - p)(\$10)$ if [H is true]...

A bet against [H] at rate $(1 - p)$: you win $p(\$10)$ if [H is false]...[Adapted from (Hacking, 2001, 2009:157)]

If your personal probability that hypothesis H was true was p , and your personal probability that H was false was anything other than $(1 - p)$, then a clever bookmaker could offer you a set of bets where you would be guaranteed to lose overall: a sure-loss (or 'Dutch book') contract.³⁶ Your set of personal probabilities would fail to conform to the probability calculus, and if the bets were real ones you would be acting irrationally in terms of personal utility.³⁷

Your personal probabilities need not stay fixed for ever though:

Unless you are truly prejudiced, new evidence should have some effect on what you believe – and on your personal probabilities. (2001, 2009:172)

³⁶ See (Hacking, 2001, 2009:164-5).

³⁷ See also p249ff.

A Bayesian explains the relation between evidence and belief in terms of conditional probability – by calculating (estimating, representing) what your personal probability for a particular hypothesis would (or should) be if that evidence were true.³⁸

(5.1) below is a formulation of Bayes' Rule. The effect of new evidence on an individual's personal probability is expressed in relation to H_j , one of the hypotheses in a 'partition'³⁹ H_i (H_1, H_2, \dots, H_n):

$$(5.1) \quad \Pr(H_j/E) = \frac{\Pr(H_j)\Pr(E/H_j)}{\sum[\Pr(H_i)\Pr(E/H_i)]}$$

Here:

$\Pr(H_j)$ is the individual's initial (prior) personal probability for the hypothesis H_j .

E is the new evidence.

$\Pr(E/H_j)$ is his probability for E , if H_j were true. (Known as the 'likelihood'⁴⁰ of hypothesis H_j in the light of evidence E .)

$\Pr(H_i)$ is his initial (prior) personal probability for each competing hypothesis H_i .

$\Pr(E/H_i)$ is his probability for E , if hypothesis H_i were true.

$\Pr(H_j/E)$ is his posterior personal probability for H_j following evidence E . This follows the 'rule of conditionalization' which states that your 'personal rate for betting on H, conditional on E, *before* you know that E is true' should equal

³⁸ This aligns with our working definition of 'evidence' from p62.

³⁹ A partition is 'a set of mutually exclusive and exhaustive hypotheses' (Hacking, 2001, 2009:172).

⁴⁰ This term, coined by English statistician RA Fisher (1890-1962), can unfortunately lead to confusion since 'probability' and 'likelihood' are synonyms in ordinary English: see Hacking (2001, 2009:167).

your 'personal rate for betting on H ... *after* you know that E is true'
(Hacking, 2001, 2009:258-9).⁴¹

E confirms H for an individual if his posterior probability exceeds his prior probability. E disconfirms H if his posterior probability is lower than his prior probability. If E confirms H , it disconfirms not- H .

Two individuals can start with very different prior probabilities ('priors') for a hypothesis, but as long as

- (i) their priors are not exactly 0 or exactly 1,⁴²
- (ii) they experience the same successive cycles of confirming (or disconfirming) evidence; and
- (iii) they generally agree on the probabilities of getting the evidence given that hypothesis ('likelihoods');

then they should converge towards the same eventual posterior probabilities. This is because the more iterations there are the less effect the difference between their priors will have on the eventual outcome.⁴³

Bayes' Rule applied to a test case

To relate the Bayesian approach to EP4 we will start with our:

(1.3) I believe you are taking recreational drugs.

Or as assertion:

⁴¹ Developed further p223ff.

⁴² Whatever the values for $\Pr(E/H_i)$ in (5.1), if $\Pr(H_j) = 1$ then $\Pr(H_j/E)$ can never be anything but 1, and if $\Pr(H_j) = 0$ then $\Pr(H_j/E)$ can never be anything but 0. Both are simple consequences of the Bayes Rule equation.

⁴³ Hacking (2001, 2009:257) for example briefly discusses this 'convergence theorem'. Also known as 'washing out the priors'.

(1.3.1) You are taking recreational drugs.

An obvious way to apply Bayes' Rule (5.1) to beliefs and assertions like (1.3) and (1.3.1) is as a prescriptive principle. Later we will question the nature of that prescription,⁴⁴ but assuming for now that (5.1) is the correct principle, we can ask: Have 'I' (the believer/asserter of (1.3)/(1.3.1)) applied (5.1) at all (implicitly or explicitly)? If so, have I applied it correctly?

The partition will be hypothesis (1.3.1) plus the 'null' hypothesis:

(1.3.1n) You are not taking recreational drugs.

Should we interpret my degrees of belief in (1.3.1) and (1.3.1n) as prior or posterior probabilities? The test case says:

(1.3.3) You are taking recreational drugs, but I have no supporting evidence.

There could be evidence supporting (1.3.1n) though. I may have little or no evidence that a person I have never met is not taking drugs – other than what if anything I know or believe about base rates in the population(s) he belongs to.⁴⁵ But if, for present purposes, we assume I know the 'you' of (1.3.1) quite well,⁴⁶ and I have no evidence you *are* taking drugs, then I would have at least some evidence you are *not* taking drugs. (I have never seen you accepting drugs but have seen you refusing them; your behaviour and appearance have never shown traits associated with drug use.)

In Bayesian epistemology 'prior' and 'posterior' probabilities have specific definitions relative to a particular piece of evidence. Prior probability is the probability before the evidence comes in. Posterior probability is the probability after the evidence comes in and (by the 'rule of conditionalization') always equals the conditional probability assuming the evidence will come in. By definition not all

⁴⁴ See p226ff.

⁴⁵ Conscious or unconscious selection of reference classes can be open to risks of eg racial or gender profiling: see p295.

⁴⁶ We will pick up this assumption later: p223.

priors can be posterior probabilities: we cannot conditionalise unless we have a prior, so we must begin somewhere. Prior and posterior probabilities only apply where credences change through Bayesian conditionalisation. (Any other way of giving a value to a credence is generally called 'fixing' a probability.)

This brings us to the contrast between **subjective** and **objective** Bayesians, who differ over the kinds of constraints applicable to priors.

All Bayesians see a prior or posterior probability as expressing an individual's uncertainty, which may vary by individual. This applies both to the belief and to any imaginary gamble it may be compared with. For example you come to an unmarked fork in a path:

If you toss a fair coin to decide which way to go, you must think that it is as probable that home is to the left, as that it is to the right. Your personal probability, for each fork, is $\frac{1}{2}$. (Hacking, 2001, 2009:152)

Here you are equally uncertain that home is to the left or that it is to the right; or that the coin will land heads or land tails. The Bayesian does not think your credence of 0.5 for heads represents an 'objective' probability any more than your credence of 0.5 that home is to the left does. Only one direction is correct. The coin will land either heads or tails. In theory you could be more certain about how it will land by knowing more about its aerodynamics and the way it was thrown. Your credence of 0.5 for heads arises from your ignorance about these details, plus your knowledge, belief or assumption that the coin and its thrower are fair.

An extreme **subjective** Bayesian (Talbot, 2013:§4.2 F(a)) would say the only rational constraint on priors is conformance to the probability calculus. In our partition of (1.3.1) and (1.3.1n) the two priors must total 1. Otherwise the agent would be exposed to sure loss ('Dutch Book') contracts.⁴⁷ Some Bayesians however like Howson and Urbach reject Dutch book arguments and regard the Bayesian 'rules of epistemic probability [as] nothing but rules of logic' (2006:51).

⁴⁷ See p214.

While the mathematics underpinning Dutch book arguments is ‘unquestionable’, they see Bayesianism as purely a ‘theory of consistent probabilistic reasoning’ (2006:301).

An extreme subjective Bayesian might even allow priors of 1 and 0 for (1.3.1) and (1.3.1n) respectively – or vice versa. Rationality would be purely a matter of probabilistic coherence, so these priors would be rational.

Less extreme subjective Bayesians might agree the values must total 1, since they exhaust all possibilities, but allow them to be any value except 1 and 0, reserving these for logical certainties and logical impossibilities respectively. Such thinkers could plausibly argue there cannot be evidence against a logical certainty or for a logical impossibility. (1.3.1) and (1.3.1n) are neither.

We should expand on ‘conformance to the probability calculus’. Part of the idealisation involved in the conception of an ideal Bayesian agent is that, if asked to estimate a prior for a particular hypothesis, the agent takes responsibility for every other related hypothesis he has a credence for, including hypotheses about how all those hypotheses connect logically. This means committing (or being prepared to commit) to quantified credences for all other hypotheses he might be entertaining which he sees as having logical connections with the hypothesis under consideration. Bayesian agents are taken to be ‘logically omniscient’.

In the case of (1.3.1) and (1.3.1n) I might for example have many other beliefs about you and the categories (reference classes) you belong to. I might consider imaginary gambles in respect of all these related hypotheses, and on pain of irrationality I should ensure the combination of odds I would accept on all these gambles – including gambles on (1.3.1) and/or (1.3.1n) – would not incur sure losses. The ‘probabilistic coherence’ constraint therefore ensures the Bayesian approach considers a believer’s entire network of related beliefs.

There may however be few or no related beliefs. The agent could in theory have a belief with few if any connections to other credences he is committed to: for example James’s ‘religious hypothesis’ (1896, 2000:214), in a partition

representing a 'genuine option that cannot by its nature be decided on intellectual grounds' (1896, 2000:205).⁴⁸ James and a subjective Bayesian would both seem to permit a range of credences between 0 (or just over 0) and 1 (or just under 1).⁴⁹

An **objective** Bayesian though would see additional rational constraints on priors, for example the Principle of Indifference (Keynes, 1921:41ff), which Strevens (1998) words as: 'in the absence of any known reason to assign two events differing probabilities, they ought to be assigned the same probability' (1998:231). Supporters of the Indifference Principle might insist that, if a ball is drawn at random from an urn filled with unknown numbers of red and black balls and no balls of any other colour, the prior personal probability of it being black can only be 0.5, because 0.5 is the only value which is 'invariant with a change in label ("red" or "black")' (Talbot, 2013:§4.2 F(b)). This is because the agent would 'have no information at all about which balls are red and which balls are black'.

The coherence of this principle can however be questioned. Bertrand's paradox returns (at least) three different answers ($\frac{1}{3}$, $\frac{1}{2}$ and $\frac{1}{4}$) to the question: what is the probability that a chord drawn randomly in a circle is longer than the side of an equilateral triangle inscribed in that same circle? There seems nothing ambiguous about what the chord of a circle is or what its length is relative to the side of a triangle, and so the principle itself appears incoherent.⁵⁰

Strevens rejects the idea of a constraint on priors based on the Indifference Principle, 'which takes us from symmetries in our *knowledge*—or more exactly, ignorance—to *rational* probabilities'. Instead he claims we actually use a 'rule of inference that takes us from knowledge of physical symmetries to knowledge of actual physical probabilities' (1998:232).

But symmetry considerations may also, like the Indifference Principle, only apply in certain scenarios, and not provide a way to ground or constrain all personal probabilities. For example it is unlikely whether symmetry would constrain the

⁴⁸ See p140ff.

⁴⁹ This scenario is considered in **Appendix A3** p285ff.

⁵⁰ See **Appendix A4** p289ff.

probability that general relativity is true or that a God exists. If the kind of principles which objective Bayesians identify or promote only apply to certain belief categories they would be little help in supporting EP4.

There are other more complex ways in which objective Bayesians have sought to fix priors but it is beyond the scope of this work to discuss or even categorise them all. Suffice it to say, it is agreed that all have problems, like Bertrand's paradox, or limitations of application like Strevens' proposal. We will therefore confine our discussion to seeing if subjective Bayesianism can help in our articulation of EP4.

Returning to (1.3.1) ('You are taking recreational drugs') and (1.3.1n) ('You are not taking recreational drugs') we might start with the idea that a (subjective) Bayesian would allow a very high value – say 0.99 – for my prior for (1.3.1), and 0.01 for (1.3.1n) – or vice versa. This is assuming the only rational constraints against these priors are that they cannot be 0 or 1 and they must total 1.

From the perspective of EP4 and its need for a criterion of justified belief in terms of sufficient evidence, we can call this the 'problem of the priors': that a Bayesian agent might consider himself – justifiably according to Bayesianism – as having sufficient evidence for any hypothesis given appropriately chosen priors. But this theoretical freedom would imply my credences for (1.3.1) and (1.3.1n) can be insulated from all my other beliefs, including any beliefs I might have about your behaviour and about drug use base rates.

A way to evade this 'problem of the priors' for the purposes of EP4 could therefore be to insist on considering those other background beliefs, despite (1.3.3) saying 'I have no supporting evidence' for belief (1.3.1). This would be in line with the duty to seek out evidence which EP4 seems to presuppose⁵¹ and which Clifford also claims.⁵²

We will therefore suggest some plausible starting assumptions, and construe 'evidence' broadly enough to include other knowledge or beliefs I have which

⁵¹ p209.

⁵² pp210–213.

might impact my credences for (1.3.1) and (1.3.1n). Perhaps originally I knew nothing about you but your age, gender and nationality. We will also assume I knew, or thought I knew, that (say) 50% of your demographic take drugs. We will take 50% as the highest possible value for (1.3.1). Anything higher could count as evidence supporting (1.3.1), which is ruled out. So I will set my priors for both (1.3.1) and (1.3.1n) at 0.5.

As we are construing 'evidence' to include constraining background beliefs it is not strictly true that 'I have no supporting evidence' (1.3.3), since I do have *some* evidence. In this test case we are therefore reading 'I have no evidence' as 'I have no evidence about you as an individual, only about base rates for the category or categories I think you belong to'.⁵³ But this only works because of the numbers involved. If I knew you belonged to a subcategory where 99% took drugs, then it might be disingenuous to claim no supporting evidence for (1.3.3), rather than saying I lacked evidence that you were not in the 1%. We have assumed I know or believe that 50% of your demographic take drugs. If I have no further evidence about you, and can treat you as randomly selected from that demographic ('fair coin'), then, assuming I am no more justified in thinking you belong to one subset of your demographic than another, then I would be no more justified in believing you take drugs ('heads') than not ('tails'). But if the 50% moves to 99% then I would be more justified in believing you take drugs than not.

My background beliefs about your age, gender and nationality, and relevant behavioural base rates in your demographic, would count as **synchronic** coherence constraints in relation to my prior of 0.5 for (1.3.1n): 'synchronic' because they apply at the same time as my prior for (1.3.1n). I would have credences in relation to all these hypotheses which are logically connected to (1.3.1n), so for the set of credences to be consistent my credence for (1.3.1n) must be a particular value, or at least within a particular range.

⁵³ I may or may not have knowledge or beliefs about relevant base rates. Here we are assuming I do, to flesh out a plausible history of Bayesian updating prior to my belief that (1.3.1).

We can then factor in E , which is all the ‘new’ evidence I have about you – now I know you.⁵⁴ Since, from (1.3.3), I have no evidence supporting my belief that (1.3.1), then what evidence I do have – including lack of evidence that you are hiding anything likely to support my belief that (1.3.1) – will support (1.3.1n). E will qualify as ‘new evidence’ in a cycle of Bayesian updating, so it will operate as a **diachronic** coherence constraint on my posterior probability: ‘diachronic’ because it applies across the updating, which comes after my original prior for (1.3.1n).

In its simple form the ‘Dutch Book’ argument only supports **synchronic** coherence, as it applies to a set of logically related credences held at the same time. If the related credences are represented as personal probabilities, and therefore as personal ‘betting rates’, the set of credences must be arithmetically consistent so as to avoid exposure to sure-loss contracts.⁵⁵ The eventual set of logically related credences will include some (including the initial prior) at time t_1 and some (including the believer’s credence(s) for the new evidence E when it exists) at a later time t_2 . Bayesian updating relies on this **diachronic** coherence, which itself rests on the ‘rule of conditionalization’ mentioned earlier.⁵⁶ a rule which, according to Hacking, could be seen as either ‘an additional maxim or axiom of the Bayesian subjective approach’, or even

*as a moral rule. One should be true to one’s former self (unless one has undergone a radical conversion in one’s life and thought) (2001, 2009:259).*⁵⁷

These requirements for synchronic and diachronic coherence are in line with our earlier suggestion⁵⁸ relating the condition of sufficient evidence to consistency between beliefs. We are also still aligned with our three-part generic analysis of moral principles⁵⁹ into (i) a moral ‘thou shalt/shalt not’; (ii) a non-moral ‘domain’ or subject matter; and (iii) a moral requirement to exercise judgment in the context of

⁵⁴ An explicit assumption: p217.

⁵⁵ But with additional assumptions the same rationale can apply to the updating of credences in response to new evidence. See eg Lewis (1999) and Teller (1973).

⁵⁶ p215.

⁵⁷ But see papers referenced in footnote 55 above for a different view.

⁵⁸ See p209.

⁵⁹ See p209 and p213.

(ii). In the case of EP4 a Bayesian might see the non-moral domain (ii) including credences expressed as personal probabilities and the axioms of the probability calculus. But a proponent of EP4 (or CP) would also see (ii) including our ‘shared epistemic asset’ or ‘heirloom’⁶⁰ This is strictly speaking non-moral, but qualifies as a context in which a ‘moral requirement to exercise judgment’ (iii) could be appropriate. The shared epistemic asset necessarily persists through time, so necessarily has a ‘diachronic’ dimension. It is hard to see how it could survive if no one in a position to protect it saw a need to be true to his former self – particularly his epistemic and doxastic self.

Returning to our iteration of Bayes’ Rule (5.1) and using its notation, we will call (1.3.1n) ‘ H_j ’ and (1.3.1) ‘ H_k ’. With a few additional assumptions, we can now estimate my posterior probability for H_j (1.3.1n).

We have already assumed:

$$\Pr(H_j) (= \text{my prior for } H_j \text{ (1.3.1n)}) = 0.5.$$

$$\Pr(H_k) (= \text{my prior for } H_k \text{ (1.3.1)}) = 0.5.$$

We now assume:

$$\Pr(E/H_j) (= \text{my probability for evidence } E, \text{ if hypothesis } H_j \text{ (1.3.1n) were true}) = 0.6:$$

As all the evidence is of you not taking drugs, my probability that this evidence would occur if you were not taking drugs would be quite high.

$$\Pr(E/H_k) (= \text{my probability for evidence } E, \text{ if hypothesis } H_k \text{ (1.3.1) were true}) = 0.2:$$

⁶⁰ See p210ff.

My probability that this evidence would occur if you were taking drugs would be lower, although not zero, as you could be a skilful deceiver. The precise values do not matter, only that $\Pr(E/H_k)$ is considerably lower than $\Pr(E/H_j)$, and that this should be plausible – following from ‘no supporting evidence’ explicit in (1.3.3).

As partition H_i is just H_j (1.3.1n) plus H_k (1.3.1), (5.1) becomes:

$$\begin{aligned}
 (5.1.1) \quad \Pr(H_j/E) &= \frac{\Pr(H_j)\Pr(E/H_j)}{\Pr(H_j)\Pr(E/H_j) + \Pr(H_k)\Pr(E/H_k)} \\
 &= (0.5 \times 0.6) \div ((0.5 \times 0.6) + (0.5 \times 0.2)) \\
 &= 0.3 \div 0.4 \\
 &= 0.75
 \end{aligned}$$

So my posterior probability for H_k (1.3.1) = $1 - 0.75 = 0.25$. If:

- (i) I continue believing you are taking drugs, with no supporting evidence and all available evidence supporting (1.3.1n);
- (ii) the assumptions we have made here are reasonable;
- (iii) 0.6 and 0.2 reflect my true personal probabilities for evidence E , assuming hypotheses H_j (1.3.1n) and H_k (1.3.1) were true respectively; and
- (iv) we presuppose the ‘rule of conditionalization’;

then I am disobeying Bayes’ Rule. My belief would be irrational according to Bayesianism. My credences would not conform to the probability calculus, and if they were interpreted as gambles which I accepted I would be exposed to sure-loss contracts.

We do not yet have an equivalent account of ‘sufficient evidence’, but this particular case does not need it. If I am to believe (1.3.1) or (1.3.1n) – rather than holding no belief either way – then I should believe (1.3.1n) which is better supported by evidence. This seems to fit our intuitions about what (and how) it is reasonable and right to believe. Applying EP4, I will be morally wrong to believe you are taking recreational drugs (1.3.1) because, given the assumed evidence, I should either withhold belief or believe you are not taking drugs (1.3.1n).⁶¹ That is all EP4 needs to reject (1.3.1). Neither CP nor EP4 insists I should believe (1.3.1n); only that I should not believe (1.3.1).

Evidence and normativity

It seems then that, as long as we assume a duty to get evidence, which would be outside Bayesianism proper, we can give a Bayesian account of why it would be wrong⁶² to believe someone is taking recreational drugs (1.3.1) without sufficient evidence. This is assuming also that it is correct to model credences as personal probabilities, to estimate the odds we would be prepared to accept for imaginary gambles based on the respective hypotheses. The prescription to use Bayes’ Rule would then derive from Dutch book arguments, such that if we did not apply it, or applied it incorrectly, we would irrationally compromise our personal utility by exposing ourselves to sure-loss contracts.

The assumption that it is correct to model credences on betting strategies derives from the first tenet of Bayesianism, which, according to Strevens (2005:1-2) is that

*...the scientist’s epistemic attitude to any scientifically significant proposition is, **or ought to be**, exhausted by the subjective probability the scientist assigns to the proposition. A subjective probability is a number between zero and one that reflects in some sense the scientist’s confidence that the proposition is true [Emphasis added].*

⁶¹ Later we will question why I would be *morally* wrong to act against my self-interest (personal utility): see p249ff.

⁶² A wrong of irrationality on the basis of personal utility, not a moral wrong: see p212 and p249ff.

He continues:

*To found its first tenet, Bayesianism must establish that it is plausible to suppose, **or reasonable to require**, that scientists have a subjective probability for every proposition that figures in their inquiry. Ramsey (1931) proposed that to have a subjective probability for a proposition is to have a certain complex disposition to act, a disposition that can be measured at least tolerably well in many cases by assessing betting behavior...*

[Emphasis added.]

The two other tenets (conformance to the probability calculus and Bayes' conditionalisation rule) are largely consequences or expansions of the first. To 'be a Bayesian' is to accept and apply all three tenets, which implies that the assumption (commitment, prescription) to 'be a Bayesian' is external to Bayesianism itself. It would be independent of any prescriptive consequences of the decision to model credences in such a way that they can be expressed as betting strategies, as it would be the prescription to do the modelling and accept the consequences.

We should now enlarge on the assumed duty to get evidence, including the duty not to ignore available evidence. Even restricting ourselves to a purely scientific context, 'available evidence' could include: the results of the experiment we are doing; and/or the results of the experiment someone else is doing next door which we know about; and/or the results of an experiment happening in another country which we know nothing about; and/or the evidence which is 'available' in the sense that we could get it from an experiment we have not yet done; and/or the theoretically 'available' evidence we could get from an experiment which is currently practically impossible; ... and so on. A decision to be Bayes rational seems to presuppose (or entail the need for) a practical approach to identifying the values to feed into Bayes' Rule.

We can only update credences in a Bayesian way if we make a reasonable effort to find evidence. If we have no evidence because we made no attempt to get any evidence, then we cannot determine posterior probabilities $[\text{Pr}(H/E)]$ without any

evidence $[E]$ to have 'likelihoods' $[\Pr(E/H)]$ about.⁶³ We must therefore decide what is 'reasonable'. For example we could claim publicly available evidence should be used wherever possible; and that care should be taken to decide what counts as evidence for what hypothesis.

These considerations divide into issues of quality and issues of will. On the quality of evidence, relating the probability of evidence given a particular hypothesis to the posterior probability of that hypothesis implies the evidence itself is completely certain (probability 1). But we may want to update our personal probability on the basis of (say) a rumour we are reluctant to discount (Hacking, 2001, 2009:182). Jeffrey's rule gives us a way to accommodate evidence which is less than certain by factoring in personal probabilities that the evidence is true and false respectively (Jeffrey, 1983; Hacking, 2001, 2009:183).⁶⁴

The question of will is more involved. In theory we can be Bayes rational without evidence, so without updating our degrees of belief. But Hacking describes Bayesianism as a model for 'rational change in belief, in the light of new evidence' (2001, 2009:171), and we should unpack what this 'rationality' could imply.

As suggested above,⁶⁵ one view sees Bayesianism as purely a 'theory of consistent probabilistic reasoning'

in which the truth, rationality, objectivity, cogency or whatever of the premises, here prior probability assignments, are exogenous considerations, just as they are in deductive logic (Howson & Urbach, 2006:301).

But our aim is to see if Bayesianism can help EP4 with an account of sufficient and/or insufficient evidence to justify belief, including the practical consequences of belief. It might therefore be appropriate (and in line with our assumed duty to get evidence) to consider an enriched Bayesianism within a context of normative

⁶³ See p215.

⁶⁴ If evidence E is only a rumour (with $\Pr(E)$ = the probability that E is true and $\Pr(\sim E)$ = the probability that E is false) then the probability of hypothesis H in the light of the rumour that E = $[\Pr(H/E) \times \Pr(E)] + [\Pr(H/\sim E) \times \Pr(\sim E)]$.

⁶⁵ See p218.

decision theory, which is ‘about how decisions should be made in order to be rational’ (Hansson, 2005:6). This appears a plausible way to flesh out the ‘rationality’ Hacking refers to. If rationality is about maximising expected personal utility then it would seem irrational to make no effort to get evidence to update our credences. If we see the rationale behind Bayesian updating as preventing exposure to sure-loss contracts – and therefore loss of personal utility – then refusing to consider or look for evidence⁶⁶ could run the risk of losing bets which do not need to be lost and/or accepting bets at inappropriate odds – which could again entail loss of personal utility. Both could thwart the assumed aim of maximising personal utility. But an endless search for unattainable conclusive evidence would not maximise personal utility either. Hence the stipulation to make a ‘reasonable’ effort. There is a cost in seeking evidence, which should be contained to a ‘reasonable’ level.⁶⁷

We could also appeal to normative decision theory to justify decisions about what the potential evidence might be for a hypothesis, what potential evidence is possible, how we might get it, how much store to set by it, and what its quality might be. Then we could know in advance what potential evidence to look for, create the conditions for, and measure. In a straightforward context the most valuable potential evidence E will be where there is a high probability of getting E if hypothesis H is true, and a low probability of getting E if H is false. If we are Bayes rational we will see that will yield a higher posterior probability for H than before we had E .⁶⁸ Such considerations suggest Bayesianism, in an assumed context of decision theory, is broadly in line with both CP and EP4. We will keep the context of normative decision theory as a plausible assumption for now, although it will need to be qualified significantly later.⁶⁹

This ‘enriched’ Bayesianism should help determine what credence to have for a hypothesis, in comparison with other hypotheses in a partition. It would certainly

⁶⁶ Which is effectively a special case of deliberate ‘partisan’ selection of evidence to support chosen credences: see p287.

⁶⁷ See footnote 69 below.

⁶⁸ Because $\Pr(H|E) > \Pr(H)$.

⁶⁹ We are touching here on Information Value Theory, which can of course be addressed more formally: for example Howard (1966). Eventually however we will be relying less on notions like this and more on a distinctively moral duty to preserve our shared epistemic asset: see p249ff.

help do this in a dynamic context, when net evidence increases or decreases over time. Considering how bad we are at statistical thinking, it can help us avoid false beliefs arising from base rate neglect: described as a ‘sharp violation of Bayes’ rule’ (Tversky & Kahneman, 1974:1124) and a ‘sin of representativeness’ (Kahneman, 2012:151).⁷⁰ It will help us move (justifiably and therefore rationally) from ‘I believe that H ’ to (say) ‘I don’t believe that H ’ or ‘I’m not so sure about H ’. But it does not yet give us a handle on when we can say ‘I believe that H ’: when we have *sufficient* evidence, if this falls short of *conclusive* evidence. This is our next hurdle. The discussion will extend over a number of sections and cover different facets of justified belief. The first real attempt to draw the threads together will be at the end of section **5.6 Indefinite intersubjectivity**.⁷¹

Evidence threshold

A first stab at a Bayesian criterion could be that evidence is sufficient when it takes our credence above a certain value and insufficient when our credence is below that value.

But this encounters the ‘problem of the priors’:⁷² that on a strictly (subjective) Bayesian approach an agent could regard himself as having sufficient evidence for any hypothesis given appropriately chosen prior probabilities. The threshold could therefore be zero.

Our counter-suggestion was that we are only interested in posterior probabilities – credences after at least some updating in relation to evidence. This would imply a duty to get and acknowledge evidence.⁷³ Our modified proposal would then be that evidence is sufficient when it takes our credence above a certain threshold

⁷⁰ For example being terrified after testing positive in a 99% reliable test for an illness affecting one in 10,000 people (Hacking, 2001, 2009:73). The odds that you have the disease will only be a hundred to one against. Casscells *et al* (1978) asked 60 Harvard medical practitioners and medical students the question ‘If a test to detect a disease whose prevalence is 1/1000 has a false positive rate of 5 per cent, what is the chance that a person found to have a positive result actually has the disease...?’ Only 18% gave the correct answer of 0.02, and 45% gave the very wrong answer of 0.95. No wonder the growth of decision-support systems implementing Bayesian algorithms in complex medical and engineering applications.

⁷¹ pp264–267.

⁷² p221.

⁷³ p226ff.

after a reasonable period of updating and insufficient if there has been little or no updating or our credence is below that threshold. We would need to explain though what counts as a reasonable amount of updating and what the threshold would be, as the following scenario attempts to illustrate.

Imagine you have a bag labelled as containing a kilogram of pure salt. It is part of a shipment which someone claims to have adulterated by substituting 10% by weight of ground glass in some of the bags. You are seriously considering two hypotheses, which make up your partition:

H_P The bag contains pure salt.

H_A The bag has been adulterated with 10% ground glass.

We will assume priors of 0.5 for each and ignore other potential hypotheses: the bag is all ground glass; the bag is 75% ground glass; the bag is full of potassium nitrate; etc.

You shake the powder and take a sample of 10 grams. You dissolve the sample in water and filter it. There is no insoluble residue. You repeat a number of times. Each test is a Bayesian update cycle where you can revise your credences. After the first result your (subjective) credence for H_P shoots up but only to, say, 0.95. You think you might have been careless in how you shook the powder and then took the sample, so if there was glass it could have all fallen to the bottom and got missed from your sample. Before the second test you mix the remaining 990 grams more thoroughly. But you get the same result: no glass. With each test result your posterior probability for hypothesis H_P gets closer to 1. But not quite 1. A diminishing but still positive chance remains that 100 grams of glass remain in the bag and by some freakish sequence of events your sampling keeps missing it.

At iteration 90 only 100 grams remain. For H_A to be true the remaining powder must all be ground glass. You test sample 91 and there is no insoluble residue. You now have conclusive evidence that H_P is true: but you did not need Bayesian updating to tell you that. So what about after iteration 90, or 80, or 40, or 10? In

theory the remaining salt could have been mixed with 100 grams of glass. But despite this, when there were still 100 grams or more of powder in the bag at what point would we say we had sufficient evidence to believe H_P ?

Although EP4 only refers to 'insufficient evidence' – to identify illegitimate beliefs – we still need a way to identify legitimate beliefs, a way to connect credence to justified full belief. If we were speaking of knowledge rather than justified belief⁷⁴ the equivalent question would be: at what point would we say we *know* H_P was true?

After iteration 90, in theory the remaining powder could be all glass (H_A) but it is overwhelmingly probable that it is all salt (H_P) – assuming these are the only live hypotheses. There seems a case for saying we know H_P is true after iteration 90, but also a case for saying that, strictly speaking, we do not know. What line we take relates to the different risk consequences of each option. The probability that the last 100 grams are pure salt would be far higher than the probability of our being right in countless everyday cases where only the most hardened induction sceptic would claim not to know. But those who say we do not know after iteration 90 might insist the last 100 grams could still be ground glass.

The view that in a case like this we can say we justifiably believe H_P at some iteration before 91 can be called a 'threshold view': that

outright belief is belief to a degree higher than some threshold value $x < 1$: one believes that p (outright) if and only if $Cr(p) \geq x$, where $Cr(\cdot)$ is one's credence function. (Clarke, 2013:3)

Clarke distinguishes this from a similar 'rational belief' variant which he calls⁷⁵ the 'Lockean view': that

one is rational to believe p outright if and only if one is rational to have a degree of belief that p higher than some threshold value y . (2013:3)

⁷⁴ See p212.

⁷⁵ Following Foley (1993).

We have already mentioned⁷⁶ one issue with this: where do we set the threshold? Another issue, rather more profound, we can illustrate by way of the ‘lottery paradox’.

Lottery paradox

The ‘lottery paradox’⁷⁷ introduces a category of cases at the cusp of outright justified belief (or knowledge). Its ‘knowledge’ version can be summarised as follows. Imagine you have bought one of a million tickets in a lottery with only one eventual winner. It is overwhelmingly likely that yours is a losing ticket. But can you say you *know* this? Surely not, because if so you can also say you know all the other 999,999 ticket holders will also lose, including the eventual winner. But you cannot know this, because you cannot know something false. The problem though is that

since there are few empirical facts you can be as certain of as that your ticket is a loser, the intuitive answer raises the worry that your empirical knowledge is exceedingly limited at best. (Douven, 2003:395)

The ‘rational belief’ version of the paradox asks whether your belief that your ticket is a loser is rational. If it is, then it seems you could rationally believe the same of every other ticket – contradicting your knowledge that the lottery has a winner. But then it seems ‘there is precious little you can rationally believe’ (Douven, 2003:395).

According to Hawthorne the problem generalises such that the lottery itself can be removed:

I am inclined to think that I know that I will be living in Syracuse for part of this summer. But ... I am not inclined to think that I know whether or not I will be one of the unlucky people who, despite being apparently healthy, suffer a

⁷⁶ p232.

⁷⁷ First described by Kyburg (1961:197-9).

*fatal heart attack in the next week. ... Indeed, I am just as unwilling to count myself as knowing about the heart attack as I am to count myself as knowing about the lottery...(2004:3)*⁷⁸

Lottery cases can cast doubt on whether we can construe 'sufficient evidence' (for either knowledge or legitimate belief) in terms of a credence threshold short of 1.

A lottery ticket holder may be more certain his ticket will lose than of his grandmother's birthday.⁷⁹ But he might say he knows his grandmother's birthday is 3 August but not that he will lose the lottery. If he knows (or has no reason to doubt) the lottery is fair, then he knows he has a finite chance of winning, which is an integral part of knowing what a lottery is.

If he did not know he had a finite chance of winning then it is debatable whether what he did would count as taking part in a lottery rather than, say, giving money away. Knowing he is taking part in a lottery, and therefore knowing he has a finite chance of winning, is incompatible with knowing he will lose, but compatible with knowing he will very probably lose. Yet even if, as a reflective individual, he is aware of his own epistemic fallibility, that is not in the same way an integral part of believing or knowing his grandmother's birthday is 3 August. He can qualify as knowing his grandmother's birthday whether or not he is aware of his epistemic fallibility.

So a lottery ticket holder puzzled by being more certain of losing than about the date of his grandmother's birthday may not be comparing like with like. Knowing he has a one-in-a-million chance of winning stops him saying he knows he will lose. But family birthdays are the kind of thing we generally say we know. We may or may not ask ourselves whether we are sure. If we did we might realise we are more certain of (say) our own birthdays than of our grandparents'. There may be more than a one-in-a-million chance we have got our grandmother's birthday wrong, and we may or may not then replace 'I know my grandmother's birthday is 3 August' with (say) 'I am almost certain my grandmother's birthday is 3 August'.

⁷⁸ This scenario is understood to derive originally from (Vogel, 1990).

⁷⁹ To use an example from Frankish (2009:79).

I would say without qualm that 'I know my birthday is 6 May'. But after reflecting on the fallibility of all empirical knowledge I might acknowledge I could be wrong. Perhaps I had been swapped at birth and my real birthday was 5 May. But this would not make me replace 'I know' with 'I think'. I would more likely stick to 'I know' while acknowledging the fallibility of all empirical knowledge. I may have good reasons for thinking the possibility of error extremely remote but not know how to be more certain. I could search in vain for evidence of a cover-up – because there was no cover-up or because it was so perfect. Even if I found apparent evidence that my birthday was 5 May how would I know this new evidence was genuine? Yet no matter how unlikely I thought it was that I was wrong about my birthday, if I thought it was possible (because of the fallibility of all empirical knowledge) and chose to give it a numerical probability, I could envisage a lottery with enough tickets that my chance of winning was even less. But my reluctance to say 'I know I will lose this lottery' would not make me retract 'I know my birthday is 6 May'. Why the difference?

The Bayesian approach allows personal credences to be modelled on imaginary gambles. To make this explicit we could construct a 'no loss' choice involving a random number generator displaying a number between 1 and n . Before the number displays I am asked to pick any ticket numbered between 1 and n , and choose between:

(5.2.1) £1,000,000 if I am right about my birthday and I give up my ticket; or

(5.2.2) £1,000,000 if the number I pick is *not* the number the generator displays.

How high must n be before I choose (5.2.2)? Imagine my credence about my birthday was only 0.9, so I would have a 10% chance of being wrong. If $n > 10$ I should choose (5.2.2), as I would have a <10% chance of picking the 'losing' number.

Now consider the case where I say I know my birthday is 6 May with no reason to doubt it – other than the fallibility of all empirical knowledge. My credence would be far higher, but not quite 1. There will be some very high value of n where I would be agonisingly conflicted between (5.2.1) and (5.2.2), and above which my choice should be (5.2.2).

So it seems an individual could have two beliefs with the same personal probability ($1 - 1/n$): one about his birthday which he would say was justified (and therefore held on sufficient evidence) and most likely also knowledge; and one about not picking the winning number which he would say was neither knowledge nor a belief held on sufficient evidence, but only a belief which was very probably true. The anomaly seems to have something to do with the policy of modelling credence on an imaginary gamble, a policy which might lead us to assume that a belief which is not about this kind of gamble can be treated as if it is. In the lottery case all the relevant possibilities are always salient but this is not so in the birthday case. There is a difference of context between the two cases.

In the case of the bag of salt⁸⁰ we were undecided whether a high credence short of 1 corresponded to 'sufficient evidence', and therefore what, short of conclusive evidence, would constitute sufficient evidence to justify believing the salt was unadulterated. There seem parallels between this case and a game of chance. In the scenario where testing has not yet reached the threshold for conclusive proof – up to and including iteration 90 – I may still have a tiny but non-zero credence that the salt was unadulterated. If I do not know I have not picked the winning ticket no matter how high n is in (5.2.2), then, strictly speaking, even at iteration 90 perhaps I cannot say I know the bag is pure salt, and therefore cannot say I have sufficient evidence to believe hypothesis H_P ? I would then only have sufficient evidence to believe H_P is very probably true. That does not seem irrational.

But the two cases are not parallel. In (5.2.2), as long as we know the value of n , we can calculate the exact probabilities of picking ($1/n$) and not picking ($1 - 1/n$) the number the random number generator displays. We know one of the n tickets

⁸⁰ See p232.

is the winner, but not whether that ticket is ours. In the case of the bag of salt we do not know whether the bag was adulterated or (if it was) how perfect the mixture was. The two cases display different uncertainties.

Hawthorne's (2004) 'Syracuse' example⁸¹ also seems like a 'lottery' case when we consider a large enough population. If (say) the lifetime risk of fatal heart attack in a healthy middle-aged male is around 2%, that means something like $0.02 \div 30$ (remaining years of life) $\div 52$ (weeks in a year) = 0.0000128 or 1 in 78,000.

So we can ask whether each member of a population of 78,000 men knows he will still be alive next week. It is highly probable that at least one does not know that, because he will not be alive. (Or consider 10 out of 780,000 or 100 out of 7.8 million. If those men cannot know, how can the others know?)

But it takes a deliberate decision to see possible contingencies in this way. Normally if you say you will be somewhere at a particular time that is taken to mean you know, or at least believe, you will be there. Only if questioned might you start adding assumptions about not suffering a cardiac arrest. If such contingencies were quantified they may come to a higher probability than winning a lottery. In gambling contexts the uncertainty – and its significance – are explicitly to the fore, but in non-gambling contexts the uncertainty is only implicit or even suppressed. This seems key.

There appears to be a tension between the ideas of 'full belief' and 'credence' where the latter is modelled on an imaginary gamble, particularly one involving transparent 'fair odds' derivable from physical symmetries.⁸² On the one hand it seems appropriate to allow a high credence just short of 1 to qualify as 'full belief'. On the other hand 'lottery' cases in particular seem to expose differences between our intuitive understanding of 'full belief' (and therefore of knowledge) and the connotations and commitments which credences just shy of 1 would entail.

⁸¹ See p233.

⁸² Not all gambles, imaginary or otherwise, involve physical symmetries. For example horse racing.

Just because a ticket-holder might never say he knows (and therefore never has sufficient evidence to believe) he will lose a lottery no matter how many tickets were issued, we cannot extrapolate from that that we would not say we know (or have sufficient evidence to believe) something just because we acknowledge sceptical epistemic doubts. On the surface lottery cases seem to show there can be no threshold value for a credence short of certainty. But lotteries are different from non-lottery cases because in a lottery all possibilities are made explicit and need to be explicitly entertained by the agent. In Hawthorne's 'Syracuse' story⁸³ we see a shift from one type of case to the other. It is a shift in the relevant possibilities we consider: a shift in the context of assessment.

This should not imply our intuitive notion of 'full belief' ignores exposure to risk. The issue seems to be in quantifying that risk. We want to see both full belief and knowledge as (sometimes if not always) 'fallible' (and therefore exposed to risk) without necessarily committing to any specific quantification of that fallibility (or risk). Perhaps more simply, we want a notion of 'fallible knowledge' which is not a contradiction in terms. We want to allow a belief to be both 'supported by sufficient evidence' and also 'fallible'.

There seems a connection between thought experiments about lotteries and heart attacks and the idea that personal certainty implies being prepared to stake any amount that your belief is true and/or die a horrible death if it is false. It is almost a *reductio ad absurdum* of the principle of seeing degrees of belief as personal probabilities, modelled on the odds we would accept for imaginary gambles.

It would take a degree of courage (if not recklessness) to stake any amount including your life on even a necessary truth. In the context of an actual bet there is always a chance you might have misunderstood something: perhaps what you thought was a necessary truth was not in fact, or a word had changed its meaning without you knowing. In this context 'horrible death' means nothing unless it is a real one. So it must be a real bet, with something real at stake, not an academic exercise to calculate probabilities. It is therefore not clear that, if

⁸³ p233.

to assign a probability of 1 to a proposition is to cease to contemplate the possibility that it is false (Frankish, 2009:79),

then

a person who gave [hypothesis] H probability 1 would be willing to accept a bet in which the person wins a penny if H is true, and dies a horrible death if H is false (Maher, 1993:133).

We started with the suggestion that our belief that p is justified as long as our credence that p – understood in Bayesian terms – is at least as high as a certain threshold⁸⁴ value, which can be less than 1. Our bag of salt discussion suggested it was at least plausible that justified belief did not require certainty. But the lottery case discussion leads to doubt whether a purely quantitative credence threshold can serve on its own as a criterion of justified belief, as we could hold two beliefs at the exact same credence value, one justified and the other not.

Where a belief could clearly have supporting evidence (medical diagnosis; drug-taking; sea-worthiness of ships; primate evolution) Bayesian modelling of credence on imaginary gambles seems to fit. So there seems little reason to deny our account of sufficient evidence for EP4 a Bayesian component. Beliefs like these will be subject to doubts relating to how complete or incomplete we think the evidence is. We can model this kind of uncertainty on the uncertainty involved when betting.

But other factors also seem relevant to what level of evidence-based credence counts as justified belief, or indeed when high credence counts as full belief. These factors seem largely pragmatic ones relating to context and risk, for example whether the agent's credence is about his chances of winning a fortune or dying a horrible death, or reflects the kinds of everyday certainty we enjoy when telling others our plans.

⁸⁴ p230ff.

A thesis like this can be labelled ‘pragmatic encroachment’, and is typically expressed in terms of knowledge rather than justified belief: ‘A difference in pragmatic circumstances can constitute a difference in knowledge’ (Ichikawa & Steup, 2001, 2017). We will see if this helps our search for a criterion of evidential sufficiency.

5.5 Pragmatic encroachment

Stanley: What is at stake

Stanley makes a case for pragmatic encroachment in respect of ‘what makes someone’s true belief a case of knowledge’ (2005:loc 19). He compares similar but not identical scenarios to assess whether we would intuitively describe them as knowledge.

His first two both start:

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks.

His ‘Low Stakes’ scenario continues:

It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long[.] ... Realizing that it isn’t very important that their paychecks are deposited right away, Hannah says, ‘I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.’

‘High Stakes’ continues differently:

Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday.

Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, 'I guess you're right. I don't know that the bank will be open tomorrow.' (Stanley, 2005:loc 118-126)

In both scenarios the focus is Hannah's self-attribution of knowledge, or lack of it. Stanley has another case of self-attribution ('Ignorant High Stakes') and two where the attributor is not the subject: 'Low Attributor–High Subject Stakes' and 'High Attributor–Low Subject Stakes'. All scenarios assume the bank will be open on Saturday; and, for our purposes, all can be recast into 'justified belief' language, by replacing 'know' with 'have sufficient evidence to believe'.

Stanley agrees with what he sees as our intuitive verdict, that only 'Low Stakes' has a correct attribution of knowledge.⁸⁵ Knowledge is therefore 'not just a matter of non-practical facts' – concerning what evidence the subject has – 'but is also a matter of *how much is at stake*',⁸⁶ which consists of '*practical facts ... about the costs of being right or wrong about one's beliefs*' (2005:loc 143).⁸⁷

If what counts as knowledge – and, for our purposes, what counts as justified belief – for an agent can depend at least partly on what is at stake for that agent, then perhaps our criterion of evidential sufficiency for EP4 could be that level of evidence sufficient to justify taking the risk of holding the belief, in view of what the agent has at stake. As what is at stake can vary from one agent to another, so could the level of evidence sufficient to justify taking the risk of holding the belief.

We have already argued for an 'enriched' Bayesianism embedded within normative decision theory.⁸⁸ Could then 'what is at stake' for an individual be that individual's maximised personal utility? Might Stanley's identification of stake differences as a criterion of knowledge (or justified versus unjustified belief) then

⁸⁵ Later we will consider the impact of a possible moral duty to prioritise a social stake over a purely personal stake: p249ff.

⁸⁶ A similar idea may have also been proposed by the Hellenistic sceptic Carneades (Thorsrud, undated:2(b)(v)).

⁸⁷ James can be seen as making a similar point in his distinction between knowing the truth and avoiding error: pp163–165.

⁸⁸ See p228.

be already implicit in how we are interpreting the Bayesian model, where personal utility can motivate the acceptance of odds in a 'gamble' on the truth value of a hypothesis? In Stanley's scenarios personal utility would then re-enter to influence whether or not the agent should see the available evidence as sufficient to justify knowledge (or full belief).

We could illustrate this with a cycle of Bayesian updating covering Low Stakes, High Stakes, and what we will call 'Fair Stakes' (not one of Stanley's):

Hypothesis H = The bank will be open Saturday morning.

$P(H)$ = Hannah's prior probability for H = (say) 0.6.

Therefore her $P(\sim H)$ = 0.4.

E = The bank was open two weeks ago on Saturday morning.

Hannah's $P(E/H)$ = (say) 0.95.

Hannah's $P(E/\sim H)$ = (say) 0.1.

Her posterior probability $P(H/E)$ calculates to 0.93, roughly 13:1.

If Hannah's priority were winning a bet ('Fair Stakes'), she would be indifferent between accepting odds of 13:1 on that the bank will be open Saturday morning or 13:1 against that it will not. But in both Low Stakes and High Stakes she wants the bank to be open on Saturday morning, not to win a bet; and in Low Stakes her concern that it might not be open on Saturday is trumped by her reluctance to spend time queuing on Friday.

'Fair Stakes' could be seen as the Bayesian paradigm, where Hannah's priority is to update her credence coherently to avoid exposure to sure-loss contracts. She could not say she *knew* the bank would be open on Saturday morning, because then she would not accept a bet at 13:1 against that it would be closed.

So although consideration of stake differences influencing attributions of knowledge or justified belief may not be implicit in Bayesianism itself, it does seem compatible with normative decision theory, which Bayesianism can be interpreted in terms of.⁸⁹ In Low Stakes it seems rational for Hannah to take her 0.93 posterior probability that the bank will be open next day as justification for her claim to full belief (or knowledge) that the bank will be open next day. But in both High Stakes and Fair Stakes it would not be rational for her to do this.

Part of what it is to be an individual (person) is to be free to be different from other individuals, to have different priorities and interests. By contrast the ideally rational Bayesian agent is constrained to be universal: in the sense of using standard rules to update his credences, and therefore being prepared to place the same bets as any other individual would who shares his credences about the relevant hypotheses. We are not contrasting 'universal' rules for updating credences with individual initial credences: that applies in standard Bayesianism. The individual differences Stanley identifies are differences in behaviour or attitude on the basis of those credences: whether and when to bet; whether and when to claim knowledge or justified belief.

Stanley's cases are where the evidence threshold appropriate for an attribution of knowledge (or outright belief or justified belief) can vary from one agent to another depending on what the agent has at stake with respect to the relevant hypothesis. We should acknowledge that this could be problematic for a principle like EP4. A criterion of evidential sufficiency varying from agent to agent might seem too subjective and therefore risk trivialising EP4. We will need to address this shortly.⁹⁰

First though we will consider an alternative take on 'pragmatic encroachment'. This is that credence itself could be influenced by pragmatic factors.

⁸⁹ See p229.

⁹⁰ p252.

Clarke: Degrees of belief vary by context

Clarke advances the view that '[d]egrees of belief change from context to context, depending on the space of alternative possibilities' (2013:1). '[O]ne's credences in a particular context' depend on the weightings one gives 'to each of the possibilities one takes seriously in that context':

offering a bet [on a belief that p] means changing the context. When the practical importance of p changes, as it must when a bet is offered, the space of salient alternatives to p may also change. When I am offered a bet on p at very long odds and/or very high stakes, I am likely to worry more about whether p is true after all and consider a wider range of alternative possibilities. (2013:9)

To be offered a life-threatening bet on what you had taken as a necessary truth changes the context from academic detachment to one where having explored a 'wider range of alternative possibilities' you might start doubting your comprehension of necessary truths in general and of this one in particular. You might then come to a non-zero (if very low) probability that you will lose the bet, without necessarily altering your conviction that necessary truths have probability 1.

As part of the idealisation involved in the Bayesian modelling of credences on imaginary gambles we assume the agent is equally willing to take either side of the bet as long as the price is right. Your credence will be a measure of the odds you would accept whichever way you were betting: p on one side; $(1 - p)$ on the other. But outside an idealisation like this the thought of being equally willing to accept, at any price, either side of 'a bet in which [you win] a penny if H is true, and [die] a horrible death if H is false' (Maher, 1993:133) makes little sense. In an idealised case a credence less than 1 for a tautology is irrational because it would expose you to sure-loss contracts. But in a real-life case it would arguably be at least as irrational, albeit for different reasons, to stake your life on any bet – including that of a tautology being true.

A 'credence-sensitivist' like Clarke acknowledges the 'tight connection' between credence and willingness to bet, but denies that 'we can easily infer one's willingness to bet in one context from one's credences in another context' (2013:10).

Allowing for change of context could explain the difference between full ('outright', 'flat-out') belief that p and being certain that p :

being certain that p seems to require stability of opinion... If one is certain that p , then it must be difficult to move one to abandon belief in p . (2013:11)

To say I am certain that p is to predict my belief will not change if my context changes – if for example I must express my belief that p to a more critical audience than expected. But this still does not commit me to having envisaged every possible change of context, including being invited to stake my life on the truth of p .

Two people may therefore hold a belief that p at credence 1 in different ways. A may believe (outright) that p at credence 1, but be prone to revise that credence to a value less than 1 if the context of his belief changes and he becomes aware of a 'wider range of alternative possibilities' than previously considered. B may be 'certain' or 'more certain' that p (also at credence 1), and be less likely to reduce his credence even despite a change in context – perhaps because his range of possibilities was wide to begin with.

This should not undermine Bayesianism. The 'standard subjective Bayesian framework ... deals with ideally rational agents' (Clarke, 2013:2), for whom 'degree of belief 1 is maximally stable' (2013:11): if one's degree of belief for a hypothesis = 1 then there is 'no way to update by conditionalisation' to get a degree of belief less than 1.⁹¹ But if, with Clarke, we admit the impact of contextual change where human agents are concerned and allow 'an agent's degrees of belief [to] vary with features of context other than her evidence', then a

⁹¹ See p216 and footnote 42.

degree of belief 1 can be influenced by a contextual factor to drop to a value less than 1. A degree of belief 1 may then not be 'maximally stable'.

We could see this flexibility as already implicit in subjective Bayesianism. We must always be able to adjust our credences for non-evidential reasons if we recognise possibilities we have excluded and now take seriously.⁹²

According to Clarke's context-sensitive view of outright belief,

*to believe that p in a context C (characterized ... by a set of possibilities) is to rule out all possibilities in C where p does not hold.*⁹³ (2013:11)

This sort of context-dependence could however be as problematic for EP4 as Stanley's 'practical facts'. The 'space of salient alternatives' is not 'evidence'. It could differ from agent to agent, and be subjective rather than objective. It might therefore seem incompatible with an evidence principle, which is another worry we will need to address.⁹⁴

Not only do they both seem equally problematic for EP4, we also have no other reason to take sides between subjective-sensitive invariantists like Stanley and contextualists like Clarke. More important is what they share: that factors other than evidence can play a role in fixing credences and/or motivating decision-making on the basis of those credences. For our purposes we can combine the ideas of 'context' and 'at stake' as the believer's 'practical interests'.

Practical interests

We have actually alluded to this idea before, in the relation between belief and action. Peirce describes the 'essence of belief' as the 'establishment of a habit'

⁹² Later we will consider a possible moral duty to include and take seriously possibilities we had previously excluded: p249ff.

⁹³ Ruling out 'all possibilities in C where p does not hold' would mean giving not-p zero probability – but only in context C.

⁹⁴ p252.

(1878, 1998:144) – ‘habit’ being short for ‘rule of action’ (1878, 1998:143).⁹⁵

Clifford maintains that to count as a belief there must be potential ‘influence upon the actions of him who holds it’ (1877, 1879:181) and that ‘no belief is real unless it guide our actions’ (1877, 1879:193).⁹⁶ For James ‘belief is measured by action’ (1896, 2000:217 footnote); and he suggests another possible link by using examples of both belief and action in his tripartite analysis of ‘genuine options’.⁹⁷

Practical interest then provides an obvious link from belief to action. You might need two kinds of reason for accepting a belief – certainly where acceptance has an element of volition. One is your reason for thinking the belief is *true*. Another is your practical interest in an outcome: you need to accept the belief to use it as a basis for action, which would be the *value* of the belief for you.

We have not yet unpacked ‘practical interest’, so we can currently interpret it in terms of ‘personal utility’ from normative decision theory.⁹⁸

For James an ‘option’ between hypotheses may be ‘trivial’ or ‘momentous’ (1896, 2000:199-200) for the believer. Talk of practical interests appears to cover the same domain, but as a spectrum rather than a dichotomy. A possible qualm with James’s indulgence of over-beliefs arising from a ‘genuine option’ is that ‘momentous’ seems arbitrary and subjective. Who decides if an option is momentous?⁹⁹ It can only be the believer – who could be capricious, histrionic, obsessional. James’s example of religious belief and Stanley’s Low Stakes scenario are almost mirror images. James permits religious over-belief if (among other things) it arises from an option which is momentous for the believer (1896, 2000:215-6). In Stanley’s Low Stakes however Hannah’s belief that the bank will be open tomorrow is legitimate (qualifying as knowledge in the original wording) because the ‘option’ is not momentous. The difference seems to be that James’s religious believer has an intense practical interest in the consequence of believing that what she wants to be the case actually is the case – that there is a God and

⁹⁵ See p102.

⁹⁶ Also for Adler ‘belief’s everyday roles [are] prominently as guides to action’ (2002:11).

⁹⁷ See p143 and p168.

⁹⁸ See p226ff.

⁹⁹ See p174ff.

therefore she could (say) have a chance of eternal life. Low Stakes Hannah on the other hand has a negligible practical interest in the truth of the belief. Someone wired differently to James's believer could see the religious option as equally momentous, but because of an intense High Stakes practical interest in the truth of the belief, rather than in the consequence of believing, any religious over-belief he had would presumably be illegitimate.¹⁰⁰

This suggests a way of addressing a scenario from **Chapter 3**: Meiland's story about the wife faced with evidence of her husband's affair.¹⁰¹ She clearly has a great deal at stake – the survival of her 15-year marriage. But her 'high stakes' seem to work like James's religious believer's to lower the required evidence threshold for believing in her husband's innocence – which, other things being equal, we could assume would be more likely to preserve the marriage than believing in his guilt – rather than, as in Stanley's 'High Stakes' scenario, to raise it:

[H]er relationship of fifteen years standing to her husband shows that evidence alone (unless it is conclusive) should not determine belief (Meiland, 1980:22).¹⁰²

So ideas of pragmatic encroachment certainly seem aligned with previous elements of our discussion. However we must remember our goal is an account of sufficient evidence suitable for EP4. Unfortunately, not only have we been unable to quantify a credence level corresponding to justified belief, it now seems we must say that, depending on the believer's context, practical interests can affect what counts as full or justified belief.

The idea that different people presented with the same evidence could legitimately acquire or hold different beliefs would generally be viewed as a voluntarist but not

¹⁰⁰ James would probably agree.

¹⁰¹ p116ff. This scenario could also qualify for 'moral encroachment' treatment: see p295.

¹⁰² The wife's 'high stakes' could equally be described as raising the required evidence threshold for believing in her husband's guilt. But this would be to focus on what she does not want, not on what she does want.

an evidentialist claim. Indeed Clifford's moral language¹⁰³ seems to presuppose at least an element of voluntarism. The more qualified wording of EP4 also allows for people holding different beliefs on the same evidence, or different people seeing the same evidence in different ways. For now we are assuming both are legitimate. A potential issue though is that letting practical interests influence what counts as sufficient evidence risks emptying EP4 of all content.

Meiland's story has an explicitly moral context, in that the wife is morally conflicted as to what to believe. Not all cases of believing have an obvious moral dimension of course. But we have consistently claimed EP4 to be a moral principle. Our next move will acknowledge this explicitly.

Whose practical interest? Whose personal utility?

Clifford condemned his ship owner for believing – however 'sincerely' (1877, 1879:178) – on insufficient evidence. The ship owner would also run aground on EP4. But now construe sufficient versus insufficient evidence in Bayesian terms and model the ship owner's relevant posterior probabilities as imaginary gambles along the lines of our test case about drug taking.¹⁰⁴ His 'sincere and comfortable conviction that his vessel was thoroughly safe and seaworthy' (1877, 1879:177) might then appear irrational if we assume accompanying high credences for the propositions that the ship 'was old, and not over-well built at the first' and that 'she had seen many seas and climes, and often had needed repairs', and testimony that 'possibly she was not seaworthy'. His combination of credences would not conform to the probability calculus, and if they were interpreted as gambles which he then accepted, he would be exposed to sure-loss contracts.

Or consider the story in terms of the 'enriched' Bayesianism embedded in normative decision theory which we have argued for previously.¹⁰⁵ If we assume fairly standard motivations, for example to minimise the risk of being wrong, then the ship owner ought to have collected some evidence – assuming the cost of

¹⁰³ See p30.

¹⁰⁴ p216.

¹⁰⁵ p228.

getting evidence did not exceed the cost of being wrong. If he collected no evidence then he was in another way irrational. Except in exceptional circumstances, even just refusing to get any evidence seems in itself irrational in terms of normative decision theory, as the assumed aim to maximise personal utility would generally be thwarted by getting no evidence whatsoever.¹⁰⁶ He could also be irrational if he has other non-epistemic interests, for example the welfare of his passengers, and does not weigh these properly.

But the ship owner might have none of these interests. He might have no interest in the welfare of his passengers, and no particular interest in not being wrong. We could then argue that he has reached his decision in a way which is perfectly rational according to both decision theory and Bayesian thinking.

So while the ship owner's belief may be irrational if we model his credences on imaginary gambles, and make a set of assumptions which might be plausible in relation to other agents, in Clifford's story he took a real gamble, and won: 'he got his insurance-money when [the ship] went down in mid-ocean' (1877, 1879:177-8). His sin was in maximising his personal utility, not in failing to. Assuming he evaded prosecution for culpable homicide and/or insurance fraud ('told no tales'), his practical interest was served by ignoring the evidence he had and not seeking any more. A metric of this kind of personal utility therefore seems inappropriate in an evidence principle like CP or EP4 expressed in moral terms.

But we can construe 'practical interest' more broadly. Clifford's overriding doxastic priority seems to be the cost of believing wrongly, in terms of the potential harm this can cause to the shared epistemic asset.¹⁰⁷ So the preservation of knowledge as a shared social asset would be what is 'at stake' for him. CP and EP4 are worded in moral terms and presuppose a social context.¹⁰⁸ In Bayesianism we can invent imaginary gambles to estimate degrees of belief. The stakes themselves are arbitrary, so could be of social rather than personal utility. For a criterion of

¹⁰⁶ See p226ff.

¹⁰⁷ See in particular **5.3 Social context of belief: 2** p210ff. Clifford's conception of the shared epistemic 'heirloom' was first introduced on p16.

¹⁰⁸ We are explicitly excluding the theoretical possibility of a moral context which is not a social context: see p212.

sufficient evidence for EP4 or CP a stake of maximised social utility – specifically the maximised value of the social asset of shared knowledge – fits better than a private jackpot. What is at stake is not at the level of the individual but at whatever level of community the individual sees himself sharing at and being a part of.¹⁰⁹

We could adapt (5.2.1–2)¹¹⁰ for example:

(5.3.1) 50% drop in global warming if I am right about my birthday and I give up my ticket; or

(5.3.2) 50% drop in global warming if the number I pick is not the number the generator displays.

This comparison gamble has consequences beyond (5.2.1–2). Although still no-loss bets, my judgments about my degree of belief now have social and moral significance.

Why should I care about global warming? But we could equally ask: why should I care about maximised personal utility? In a sense the second question answers itself, as my personal utility is by definition what I care about. The utility values in the imaginary gamble have no pre-assigned content. They are just quantified items which the Bayesian agent happens to value. It may be more intuitive to construe credence updates in terms of potential debits and credits in a personal bank account. But the Bayesian agent need not be the account holder. He just needs to value any gain or loss. And a decision to prioritise personal gain has no less moral significance than a decision to prioritise protecting our shared environment or our shared epistemic asset.

This should not back us into circularity. EP4's moral requirement for sufficient evidence to support belief does not force us to stipulate the criterion in completely non-moral terms. It would be circular if the criterion was 'a level of evidence sufficient for the belief to avoid the charge of being [morally] wrong'. But we are

¹⁰⁹ See Singer (2011), which expands on *οἰκείωσις*, a theme of Hierocles and other Stoics.

¹¹⁰ See p235.

not saying that. We are still trying to construe ‘sufficient’ and ‘insufficient’ evidence in generally Bayesian terms, but wanting to see the utility ‘currency’ as something shared and social rather than merely personal. If that implies a moral judgment that the Bayesian agent must prioritise shared social utility, so be it. A currency of purely individual utility equally implies a moral judgment that the Bayesian agent may prioritise that.

So a stake of maximised social utility – particularly in relation to the shared epistemic asset – seems at least permissible as a currency for the imaginary gambles which can be used to quantify credences in Bayesian updating. It would also make sense as a potential factor when an individual decides (or acknowledges) his ‘context’ – and therefore his ‘range of alternative possibilities’ (Clarke, 2013:9).¹¹¹ If my context can change when invited to stake my life on the truth of a belief, it can change when seeing my potential over-belief contributing to ‘pestilence’, ‘plague’ or ‘savagery’ (Clifford, 1877, 1879:184-6). We might also rethink the attribution of knowledge or justified belief in Stanley’s ‘Low Stakes’ scenario (2005:loc 119)¹¹² – which presupposes we only need consider a stake of personal utility.

We said earlier that practical interest in the sense of either ‘context’ (Clarke)¹¹³ or ‘what is at stake’ (Stanley)¹¹⁴ introduces non-evidential, and possibly subjective rather than objective, considerations into legitimate belief. This risked trivialising a principle which should be grounded in ‘objective’ evidence. But we are now proposing a normative (and universalised) perspective by introducing the practical interest we ought to have, which is to preserve, protect and enhance our shared epistemic asset. The ‘sufficient evidence’ which we need as a criterion for justified belief might then manage to escape the charge of ‘subjectivity’.

We have not yet said much to justify the idea that we ought to have this shared practical interest, so this is what the next section will attempt.

¹¹¹ p244.

¹¹² p240ff.

¹¹³ See p246.

¹¹⁴ See p240.

5.6 Indefinite intersubjectivity

The normative notion of 'the practical interest we ought to have' is essentially Clifford's: our duty to preserve, protect and enhance our shared epistemic asset. This asset would at any moment be a compound of knowledge and true beliefs, plus credences, guidelines and so on; and although EP4 talks of belief rather than knowledge we will again find it helpful to articulate some of our ideas by considering knowledge as well as belief.¹¹⁵ Someone obeying EP4 would want to ensure his beliefs are as justified as they would need to be to qualify as knowledge, and avoid beliefs whenever possible which do not count as knowledge. In general, he would want to have the true belief that p because he wants to be in the position of knowing that p .¹¹⁶

Shared facts and the ethics of belief

Comparing 'knowledge' language with 'belief' language suggests a link between the concept of knowledge (versus mere true belief) and the kind of context where truths need to be shared. This is not to imply true beliefs are never shared, or that this never matters. But there is a sense in which the full value of shared true beliefs is only realised when they count as knowledge.

'I believe that p ' and 'S believes that p ' relate to the individual and his attitude towards actual or possible states of affairs which may or may not be true. Beliefs about actual or possible states of affairs may refer to a truth shared with others. But if 'know' is factive¹¹⁷ then 'I know that p ' and 'S knows that p ' make direct claims about that shared truth. S_1 saying 'I believe that p ' and S_2 who believes that p may be sharing the same attitude to proposition p : they are both holding p to be

¹¹⁵ See p212.

¹¹⁶ Although the focus here is on shared beliefs and knowledge in the context of a shared epistemic asset, it appears at least congruent with approaches which take knowledge (rather than just truth) to be the aim of belief, for example 'since the claim of full belief is to the truth of its content, a belief satisfies that claim only if it amounts to knowledge' (Adler, 2002:7) and 'knowledge is the standard against which beliefs are compared to see whether they count as justified' (Bird, 2007:84). See also Williamson (2000) and McHugh (2011).

¹¹⁷ The claim that 'know' is factive has not gone unchallenged: eg (Hazlett, 2010). But we are talking in the context of a community's interpersonal interaction with its shared stock of knowledge and true beliefs, where it seems safe to assume the general usage of 'knowledge' language is such that 'know' is seen as factive.

true. But the truth of 'I believe that p ' or ' S_2 believes that p ' does not entail the truth of p as do the truth of 'I know that p ' and ' S_2 knows that p '.¹¹⁸ And if p is true then p is a truth shared between S_1 and S_2 , and shareable with others.

This might be purely an implication of the factiveness of knowledge, but it is an important one to draw out. Of course even a false shared belief is shareable. But a false shared belief is not a shared *truth*. The factiveness of knowledge guarantees that what is shared is shared truth. Two people holding the same true belief are sharing the same proposition, which happens to be true, but they are not sharing it in the knowledge that it is a truth.

This may be a tiny facet of Clifford's shared epistemic asset – which in his eyes has social and interpersonal importance, both in its composite utility value and in how it helps to 'bind men together, and to strengthen and direct their common action' (1877, 1879:183). But it is a tiny step in the right direction, which we have not made while we keep belief isolated from knowledge. Clifford's ethics of belief takes as axiomatic a shared goal to acquire and enhance knowledge as a shared asset. True belief is part of that shared asset, but its value is only fully realised when it qualifies as knowledge. It is only when a shared true belief counts as knowledge that it is shared in the knowledge that it is truth. Otherwise it is just the belief that it is true which is shared.

An important justification for a moral dimension to how we acquire, maintain, confirm, modify and reject our beliefs is the roles these play in how our shared knowledge base is created and enhanced. Clifford sees us having a shared duty to protect and develop this asset, similar to our shared duty to protect and develop the bonds of trust holding our communities together. Essentially the same assumption underpins EP4, from which perspective one who shares the benefits of the epistemic asset without sharing custodianship is a kind of cheat.¹¹⁹

This leads naturally into the concept of an epistemic community, that community which contributes to, and enjoys the benefits of, the epistemic asset which its

¹¹⁸ See p264.

¹¹⁹ See also **Appendix A3** p288.

members share. We can see this community from at least two perspectives. We could see it as an already established 'club' which individuals can choose whether or not to join. An individual could do an evaluation to see if the epistemic duties of membership are worth the epistemic benefits. An honest and morally blameless individual could decide not to join, as long as he does not 'cheat' by enjoying benefits without reciprocating any duties.

The ship owner would be one such cheat, assuming he made at least some use of the knowledge and skills of whatever communities he belongs to. It would be hard to imagine him surviving in business without doing so. He is therefore immoral at least twice over. He is criminally negligent of his passengers' safety, and profits from their misfortune. But he is also an epistemic cheat: not helping preserve and protect the shared epistemic asset from which he benefits.

But we can also see an epistemic community from another perspective, from which the ship owner is not just a cheat but his doxastic negligence endangers the very concept of knowledge itself. From this perspective we could argue that without the shared practical interest there can be no epistemic community and therefore no knowledge or belief at all.

To flesh this out we will call on Edward Craig's 'state-of-nature theory' (1990, 1999:10) of knowledge. Belief may not aspire to be knowledge by definition, but Clifford is saying (and EP4 agrees) that we ought to believe in this way. Which suggests that an account of evidential sufficiency for knowledge might help us towards an account of evidential sufficiency for belief suitable for EP4.

Craig's 'good informant'

Craig explores what he sees as the conditions and presuppositions for knowledge behaviour and knowledge language in the context of the kind of community where such behaviour and language can take hold, including

what the concept of knowledge does for us, what its role in our life might be, and ... what a concept having that role would be like... (1990, 1999:2)

His initial summary is that the concept of knowledge is used to 'flag approved sources of information' (1990, 1999:11). Because

[h]uman beings need true beliefs ... to guide their actions to a successful outcome ... they need sources of information that will lead them to believe truths. (1990, 1999:11)

Humans 'act as informants for each other', so have 'an interest in evaluating sources of information'.

'Informant' is a key role for Craig. Its complement is 'inquirer': 'someone seeking information on the point whether or not p ' (1990, 1999:12). The inquirer will want an informant with respect to p such that both:

Either p and he believes that p , or not- p and he believes that not- p

and

If he tells us that p , we shall thereupon believe that p . (1990, 1999:13)

This is because 'we do not just want to have truths enunciated in our presence' but 'to be brought to believe them'.

Craig starts not with human needs but with a rather more simple and 'individualist' creature's. A barnacle for example has 'a certain need' so wants 'something which, there and then, will satisfy the need' – food, say – while having 'no cause to distinguish ... the presence of food, its own capacity to be nourished by that food, [and] its own capacity to detect the food and reach it' (1990, 1999:82-3).

Moving up the intelligence scale we see this 'primitive holism' start to 'fragment' (1990, 1999:83), with the ability to distinguish, say, 'food here, soon' from 'food there, soon': different situations calling for 'different strategies'. Increasing socialisation and language use add further layers of sophistication. Others may be

'better ... at recognising food' or 'may be able to use substances as food which it cannot'. The picture develops of

natural pressures driving thought away from the totally subjectivist stance, the pure here and now for me as I am here and now ... [towards ideas of] the satisfaction of a need at other times and other places ... [and of] recognitional and behavioural (and maybe digestive) capabilities other than those I have here and now, and hence of an object which can in the right circumstances satisfy such need whilst ... coming nowhere near to meeting the wholly subjectivised conditions from which our thought experiment began. (1990, 1999:83-4)

A critic might suspect an unwarranted 'just-so story'. But Craig is not offering an evolutionary theory of how the concept of knowledge actually developed. His claim is instead that

the core of the concept of knowledge is an outcome of certain very general facts about the human situation[.] ... Given those facts, and a modicum of self-conscious awareness, the concept will appear; and for the same reasons as caused it to appear, it will then stay. (1990, 1999:10)

Among these facts are that humans are living, sentient, intelligent and intercommunicating social beings whose survival depends on the satisfaction of needs which can be met in a variety of ways at different levels of cost and risk, and whose shared languages incorporate shared concepts which have arisen in response to necessity.

Craig's next move is to 'explain why we have objectivised concepts' (1990, 1999:84). He sees 'objectivisation' (1990, 1999:82) as a 'general principle' which is 'widely involved in concept-formation'. His initial example is the concept of a chair, which starts from a purely subjective interest in 'something I can now sit on' (1990, 1999:84). In beings such as ourselves this will lead to an interest in 'objects I could sit on if [or when] I wanted to' and then in 'hearing the opinions of others as to where there are objects which I can sit on if I want to'; and eventually the

concept of something which is, in abstraction from what any particular person wants at any particular time or place [etc] simply suitable for sitting on

– all the while moving towards greater and greater ‘objectivisation’. Craig claims his account

need not presuppose that the wholly egocentric, ‘subjectivised’ thought from which it began actually exists or existed ... only that if it exists, at any time, or in any individual, it will develop in the direction of objectivisation.

He next applies this thinking to the ‘situation of the inquirer and the concept of the good informant’, again starting at the ‘most subjective’ extreme. If I am seeking information as to whether or not p , I will want an informant who is satisfactory for my purposes, ‘here and now, with my present beliefs and capacities for receiving information’: someone who is

- (1) accessible to me;
- (2) recognisable by me as likely to be right about p – in the sense of possessing a ‘property which correlates well with being right about p ’ (1990, 1999:85);
- (3) ‘as likely to be right about p as my concerns require’; and
- (4) able to communicate with me.

Craig then considers various ways in which these four conditions can fail, still viewed from the ‘subjective’ perspective of an individual inquirer. In respect of condition (1) the informant may not be immediately accessible, or not accessible at all. Different ways of failing condition (2) may depend on the inquirer’s specific attributes of, say, ‘sensory acuity, intellectual ability, theoretical background; [or] spatial position or bodily orientation’. So whether an inquirer (subjectively) rates

someone a good informant can depend on many features of the inquirer's situation and his relationship with the informant (1990, 1999:87).

Now consider a range of different inquirers: all with 'the same problem – how to come by the truth whether p ', but with different ways of trying to solve it depending on their individual needs and circumstances. Most importantly

these individuals form a community, and are in some degree at least helpful to others and responsive to their needs. And even if I, as one of the community, am not so inclined, I shall still need an appreciation of their point of view if I am to be any good at getting them to help me. From such facts arise a pressure towards the formation of 'objectivised' concepts, concepts which separate ... the common core from the multitude of accretions due to particular persons, and so varying with them. (1990, 1999:87-8)

I will for example hope 'that others will on occasion recommend informants to me' (1990, 1999:88) – because there will be others who, in respect of condition (2),

can detect properties of the informant which I cannot detect, ... or have more knowledge than I have of which properties correlate well with being right on the topic at issue.

As objectivisation proceeds the requirements to be met by a good informant fare differently. The accessibility (1), detectability (2) and communication (4) conditions will diminish as potential attributes of an objectivised concept of a good informant. Not because they no longer matter to the inquirer, but they qualify not as part of the objectivised concept of a good informant but as logistical features of each individual inquirer's context, and will differ depending on the circumstances and attributes of each inquirer. Both the number and variety of ways in which they can be met will increase as objectivisation proceeds.

Increasing objectivisation will also affect condition (3), about the informant being 'as likely to be right about p as my concerns require' (1990, 1999:85). But removing individual inquirers' characteristics makes condition (3) more rather than

less stringent, edging us 'towards the idea of ... a good informant as to whether p whatever the particular circumstances of the inquirer' (1990, 1999:91): for example how important it is for him to get to the truth, and what is his attitude to risk. This condition therefore differs from the other three by relating to the content and quality of what the informant can impart, not the logistical details of the inquirer's circumstances.

To illustrate Craig's 'principle of objectivisation' in action we will offer an example which is not one of his and which deliberately avoids any 'historical' or 'developmental' aspect. It does however follow Craig by starting at the 'most subjective' extreme.

Imagine I am a neurologist and S_1 is my patient or experimental subject. I ask S_1 to hold ice in his right hand while I fill two buckets A and B from the cold tap. I ask S_1 to drop the ice, put his right hand in bucket A and his left in B, and then tell me if the water is the same temperature in both. S_1 says:

(p) The water in A is warmer than the water in B.

In this scenario S_1 would be as 'right about p as my concerns require' – because my (very specific, therefore purely 'subjective') concern as a neurologist is to test S_1 for normal sensory function. Had my concerns been different, say I had wanted to know if both buckets had been filled at the same time from the same tap, S_1 's statement that p would have been uninformative: it is indeed false. In those circumstances S_1 would not be a good informant.

Moving towards greater 'objectivisation' we now expand the number and range of possible inquirers. S_1 would not count as a good informant in the circumstances of a range of actual or potential inquirers about whom we might have no idea why they would be interested in S_1 's response to the question, and no idea what else they know or believe which might lead them to be informed or misled by what S_1 says.

To meet more 'objectivised' criteria another informant (S_2) might answer with the

water temperatures of each bucket. Or he could refuse to hold ice in his right hand before putting it in the water; or he could change p to 'The water in A *feels* warmer than the water in B'; or qualify it with '...but I had been holding ice in my right hand'. An S_2 able to inform an indeterminate range of people in a broader and ultimately indefinite variety of circumstances about how the two buckets compared would be a more universally valuable informant than the original S_1 . He would be one the community would be more likely to recommend as a 'good informant'.

Here we are not concerned with whether S_2 gained the belief by means of a reliable process, although he very probably did. Our focus is not on how S_2 came by his belief but on the range of actual or potential inquirers likely to be interested in what S_2 could inform them of.

We can interpret the story in terms of 'sufficient' and 'insufficient' evidence. The original S_1 has some evidence (the water in A feels warmer than the water in B) but this is insufficient for us to say either that S_1 *knows* the water in A is warmer than the water in B or that S_1 's belief that the water in A is warmer than the water in B is justified. A 'better' informant like S_2 might dip a thermometer in each bucket and report that the temperature is the same: we would probably say S_2 does know – and therefore that his evidence is sufficient to justify his (true) belief – that the water in both buckets is the same temperature.

According to Craig's account we identify one who 'knows that p ' as one who not only holds the true belief that p but is also someone we would recommend – and who 'recommends himself' (1990, 1999:88) – as 'a good informant as to whether p whatever the particular circumstances of the inquirer' (1990, 1999:91).¹²⁰

Craig's account sheds useful light on scepticism. The further we move towards greater 'objectivisation' the less we can assume about the nature, capacities and needs of the potential inquirer. In my own individual (therefore 'subjective') case my informant needs to be 'as likely to be right about p as my concerns require'. But the further we move from 'my concerns' to an increasingly objective 'anyone's

¹²⁰ For possible relevance to 'moral encroachment' cases see p295.

concerns' the further we will

*shift towards adopting a high value for the likelihood [to be right about p] required before we are willing to recommend an informant. For the conditions under which we recommend informants will not, in general, be ones in which the recommender is sufficiently aware of the concerns for which the informant and his information are needed. This fact will push the required standard up to such a level that the recommender may responsibly issue his recommendation whilst knowing nothing of them, that is to say, to a level that the recommender may reasonably take to be high enough to satisfy **all, or all practical**, purposes. Our use of 'knows' ... marks the attainment of that level. (Craig, 1990, 1999:98). [Emphasis added.]*

A radical sceptic would insist on certainty so the level would need to satisfy 'all' purposes. In theory a potential inquirer could be a Cartesian doubter demanding an informant who would still be right about p even if an evil genius were manipulating all relevant sensory experiences. A 'fully-objectivised-but-still-practical' stance on the other hand would settle for a level high enough to satisfy 'all practical' purposes:

In everyday practice we happily bandy the word 'know' about without having to feel that our chances of being wrong are literally zero... (1990, 1999:102)

We therefore see two potential vanishing points at the end of Craig's 'road of objectivisation' (1990, 1999:91) leading to knowledge. There is the theoretical extreme of 'absolute certainty' conceived to satisfy a Cartesian sceptic; and the fully-objectivised-but-still-practical extreme. We do not have to see the latter as imperfect or second best, even if its context is where the probability of a belief being true might be less than 1. We will seek a less quantitative way to delineate these two extremes.

The sceptical stance leads to something like solipsism. But Craig's account evades this, by committing us to what we might call 'intersubjectivity': a community of interacting individuals with at least some shared interests and values, and who

acknowledge and assume each other's subjectivity. Otherwise there is nothing to motivate the demand for a 'good informant'. At its most objectivised extreme the demand could be for an informant 'good' enough to satisfy a Cartesian doubter. But this pushes us into a contradiction. We must imagine a context where a 'recommender may responsibly issue his recommendation whilst knowing nothing' about the 'concerns' of an indeterminate set of potential inquirers – including a Cartesian sceptic. But a sceptic could doubt the reality of both the informant and the context of recommendation and demand which has summoned the sceptic into existence to serve as part of the criterion for the informant's possession of knowledge.

Craig's account does not claim to disprove scepticism. But it sidesteps it by positioning the sceptical stance as inappropriate. We cannot follow his 'road of objectivisation' to the vanishing point of theoretical scepticism – because his road cannot enter a region which forbids intersubjectivity.

Adopting Craig's account therefore commits us to intersubjectivity. But EP4 itself also presupposes intersubjectivity. Earlier we mentioned the 'shared dimension' to justified belief in a social and intersubjective context, such that the 'true value of shared true beliefs is only realised when they count as knowledge'.¹²¹ This is because Clifford's ethics of belief, the spirit of which EP4 shares, assumes a shared goal to acquire and enhance knowledge as a shared asset, and key to this is how we acquire, maintain, confirm, modify and reject our beliefs.

To guarantee that S's true belief that *p* qualifies as part of our shared epistemic asset – that is, as knowledge – 'the reason(s) S has for believing that *p* would need to be valid for others' who share that asset. This does not commit us to specifying what the justifying mechanism might be which would make those reasons 'valid for others'. Following Craig's 'principle of objectivisation', we can expand those 'others' into an indefinite range of others, with an indefinite range of concerns, constituting an epistemic community of inter-communicating conscious beings. Our second, 'practical', vanishing point becomes an extreme of indefinite

¹²¹ p253.

intersubjectivity.

This chapter's objective was to get to a criterion of sufficient evidence for justified belief. So far we have arrived at Craig's criterion for something (proposition, belief, utterance) to count as *knowledge*. An informant would count as knowing that p if he not only holds the true belief that p but is also someone we would recommend, and who 'recommends himself' (1990, 1999:88) – for 'all practical purposes' (1990, 1999:98) – as 'a good informant as to whether p whatever the particular circumstances of the inquirer' (1990, 1999:91). The level of evidential sufficiency for knowledge would be the level at and above which this would be true. Can the same criterion apply to justified belief?

A criterion of sufficient evidence for EP4

A key difference between belief and knowledge is that knowledge implies truth whereas belief does not. Someone making a first-person knowledge claim is implicitly making a first-person claim to truth:

'I know that p ' $\supset p$

Someone making a third-person knowledge claim is also implicitly making a first-person claim to truth:¹²²

'S knows that p ' $\supset p$

But there is no implicit truth claim in a third-person belief claim:

'S believes that p ' $\nrightarrow p$

Whether there is an implicit truth claim in a first-person belief claim will depend on how we interpret that first-person belief claim. 'I believe that p ' can be true while p is false. So to that extent

¹²² However a third-person report of a first-person knowledge claim would not imply a first-person claim to truth: 'S says, "I know that p "' $\nrightarrow p$.

'I believe that p ' \nrightarrow p

Even though from a first-person perspective I cannot claim both that 'I believe that p ' and ' p is false' without committing Moore's paradox,¹²³ knowledge and justified belief still differ in that knowledge implies truth whereas even justified belief does not. This is despite the fact that 'justified' for us now implies a justification sufficient for the belief to qualify for inclusion in the shared epistemic asset. We are not saying 'I justifiably believe that p ' is equivalent to 'I know that p ': no matter how justified I am in believing that p , p could still be false, whereas if I know that p then p cannot be false.

Craig's 'state-of-nature' theory gives us a criterion of 'sufficient evidence' for knowledge:

A level of evidence sufficient to support the level of justification required to be 'a good informant ... whatever the particular circumstances of the inquirer' (1990, 1999:91)

If the evidential justification for belief is such that the belief qualifies for inclusion in the shared epistemic asset – that is as knowledge – then, in theory at least, the same criterion should apply to justified belief.

We can ask why anyone would want to be an indefinitely 'good informant', why anyone should value, promote, facilitate and generally be a part of that 'intersubjectivity'. This question should not arise for Craig, for whom the idea of an indefinitely 'good informant' is what explains the objectivised concept of knowledge. This is not in itself a moral concept – which is not to say he sees his state-of-nature account of knowledge devoid of moral implications:

[T]he two facts that human life is social, and that the members of a human society have differing capacities and perspectives, make it obligatory to form

¹²³ See p28.

the separate conceptions of the state of the object [in this context, knowledge], which is invariant with respect to different individuals, and the states of the 'consumers' of the object, which are not. We are in the same area here as are those ethical theorists who point out the social advantages of adopting moral principles which prescind from facts specific to individuals. (1990, 1999:90)¹²⁴

Our focus though is belief, so our question is: why should anyone want his beliefs to be justified sufficiently so as to qualify for inclusion in the shared epistemic asset? We would argue it is that same shared objective of social utility – and specifically the maximised value of the social asset of shared knowledge – which EP4 presupposes.

An immediate potential objection is that our proposed criterion appears too stringent to apply to every belief. At times we may reasonably expect an agent to have this level of evidential justification, but can it apply across the board? Our response is to remember once again that EP4 is a moral principle and like all moral principles requires the exercise of judgment.

It is not hard to invent scenarios where it would be virtually impossible to avoid, say, 'stealing' or 'harming an innocent person' without detailed context-specific knowledge few agents could be expected to possess: how were you to know a child had deliberately left that coin on the pavement as part of a game, or that the person you went to shake hands with suffered from acute haphophobia? The criterion must therefore be seen as having an aspirational, 'wherever possible', element – appropriate in the context of a moral principle like EP4. We would be guided in our judgment by our shared practical interest in the preservation of our shared epistemic asset because, if Craig is right, without that shared practical interest we might have no epistemic community and therefore no knowledge or belief at all.

¹²⁴ This also seems to be a similar area to the moral meta-principle of 'universality': p188.

So we should also be able to address another previous complaint,¹²⁵ that practical interest introduces non-evidential, and possibly subjective, considerations into legitimate belief – ostensibly at odds with a principle based on a criterion of sufficient evidence. Practical interest may explain why we actually do adjust our credences depending on the ‘space of salient alternatives’ (Clarke, 2013:9) our context makes us take seriously,¹²⁶ or why our everyday attributions of knowledge (or justified belief) do sometimes reflect ‘how much is at stake’ (Stanley, 2005:loc 143) for ourselves and others.¹²⁷ But EP4 states that, subject to caveats and qualifications, it is morally wrong to acquire or hold a descriptive belief without sufficient evidence, regardless of what we, with all our imperfections, actually find ourselves doing. So a hard line would be to say that practical interest must fall away as being too agent-relative to play any part in EP4.

But an alternative line would be that all that has to fall away is the agent relativity, as this is what provides the variable component. The shared objective of social utility – and specifically the promotion and protection of our shared epistemic asset – is now the ‘practical interest’ we ought to have and which ought to guide our epistemic behaviour. This seems plausible both as the context we ought to see ourselves in if we are to obey EP4, and the stake we should not risk by disobeying it. In practical terms the ‘hard’ and ‘alternative’ lines come to the same thing, as both remove the individual agent as a determinant of what qualifies as sufficient evidence, replaced by a universalised, socially-oriented imperative which the individual agent must measure up against.

Concluding summary

This chapter’s goal was a criterion of evidential sufficiency suitable for EP4,¹²⁸ which we have now reached. This is:

¹²⁵ See p252.

¹²⁶ See p244ff.

¹²⁷ See p240ff.

¹²⁸ p206.

Wherever possible, a level of evidence sufficient to support the level of justification required to be ‘a good informant ... whatever the particular circumstances of the inquirer’ (Craig, 1990, 1999:91)

In terms of our three-part analysis¹²⁹ of moral principles in general, this addresses both the (not completely non-moral) ‘domain’ (ii) and the moral requirement to exercise judgment (iii) in respect of that domain.

The concept of the shared epistemic asset¹³⁰ underpinning both CP and EP4 ends up doing much of the heavy lifting in relation to our eventual criterion. But first we tried to see how far we could get without it, considering Bayesian¹³¹ accounts of the relation between evidence and belief.

Bayesianism sees degree of belief as a personal probability which can be modelled as the betting rate an individual would accept as fair in an imaginary gamble.¹³² An initial challenge was the ‘problem of the priors’,¹³³ to which our response was to consider only credences after at least one update cycle. We therefore assumed a duty to get evidence and to take that evidence seriously.¹³⁴ This made sense if we adopted an ‘enriched’ Bayesianism considered within normative decision theory,¹³⁵ which also explained both the normative aspect of Bayesianism (‘why be a Bayesian?’) and the currency of personal utility which Bayesian imaginary gambles could be denominated in.

Although evidential sufficiency and insufficiency are not strictly part of Bayesianism, this does not make Bayesianism irrelevant to EP4. Bayesianism goes a long way towards explaining ‘good practice’ in our doxastic behaviour in support of decision-making under risk, particularly in respect of the relation between degrees of belief and evidential support and the avoidance of doxastic fallacies – even to the extent of outsourcing decision-making to mechanised

¹²⁹ p208ff.

¹³⁰ p210ff.

¹³¹ p213ff.

¹³² p214.

¹³³ p221.

¹³⁴ p221.

¹³⁵ p228.

Bayesian algorithms in crucial medical or engineering contexts.¹³⁶ It also suggests an interpretation of *insufficient* evidence applicable in some contexts, which would be where an agent's posterior probability for a hypothesis following at least one update cycle is less than that of its negation.¹³⁷

We need something much stronger than this however: a criterion of sufficient evidence short of conclusive evidence, as conclusive evidence needs no Bayesian explanation.¹³⁸

We considered the idea of a threshold.¹³⁹ But apart from the obvious issue of deciding and agreeing its numerical level, the mere concept of a threshold seems unworkable on its own. 'Lottery' cases¹⁴⁰ show we could have two beliefs at the exact same personal probability, one we would say we justifiably believe (or even know), and the other we would say is only very probably true.

To address this we considered two different approaches to 'pragmatic encroachment'.¹⁴¹ Both appeared to explain facets of everyday epistemic language, and also aligned with our assumed 'enriched' Bayesianism embedded in normative decision theory: the 'practical interest' motivating 'pragmatic encroachment' seeming equivalent to the 'personal utility' representing the stakes of the 'imaginary gambles' which Bayesian credences can be illustrated as.

But a conception of practical interest restricted to personal utility seems inappropriately agent-relative for an EP4 expressed in moral terms. This is particularly evident when relating our developing notion of evidential sufficiency to, for example, Clifford's ship owner story.¹⁴² From the perspective of personal utility, the ship owner's interest was served by believing as he did, not by withholding belief on insufficient evidence.

¹³⁶ See p230 footnote 70.

¹³⁷ p224ff.

¹³⁸ p231.

¹³⁹ p230ff.

¹⁴⁰ p233ff.

¹⁴¹ p240ff.

¹⁴² p249ff.

For the purposes of EP4 therefore there seems every reason to invite back into our account the dual conception of our shared epistemic asset and our shared responsibilities to that asset. This transforms the 'practical interest' in pragmatic encroachment into the 'practical interest we ought to have', which then provides an explanation for our duty to get evidence and take evidence seriously, and for the 'utility' currency denominating the stakes of Bayesian imaginary gambles.

It remained to make some inroad into justifying why we should have this shared practical interest (without trying to derive an 'ought' from an 'is'¹⁴³) and then to factor in a criterion of evidential sufficiency for belief. For both we enlisted the help of Craig's 'state-of-nature' theory of knowledge.¹⁴⁴ Although the shared epistemic asset accommodates both true belief and prescriptive principles for its own onward custodianship, its core is knowledge. So key to how we honour our duty to preserve, protect and develop our shared asset is how we aim to ensure our beliefs could qualify as knowledge.

Craig offers a definition of knowledge in terms of the requirements of a 'good informant ... whatever the particular circumstances of the inquirer' (1990, 1999:91) as an explanation of how the concept of knowledge arose, not as a suggestion of what (morally) ought to be the case. With EP4 however we are in explicitly normative territory: the 'ought' in respect of belief is that our doxastic behaviour should be such that our beliefs qualify for inclusion in the shared epistemic asset. It is not that our beliefs must be identical to knowledge – knowledge and belief cannot be the same thing – but that our beliefs should be as justified as our knowledge. Our eventual criterion¹⁴⁵ is therefore:

Wherever possible, a level of evidence sufficient to support the level of justification required to be 'a good informant ... whatever the particular circumstances of the inquirer' (Craig, 1990, 1999:91)

¹⁴³ See for example pp62–72.

¹⁴⁴ p255ff.

¹⁴⁵ p264ff.

The imperative to be a 'good informant' commits us to intersubjectivity, allowing us to navigate around radical scepticism and avoid consigning virtually all our beliefs to a black hole of solipsism.

We mentioned¹⁴⁶ the risk that our search for a criterion of sufficient evidence could end in a trivial and/or circular formula like 'a level of evidence sufficient for the belief to avoid the charge of being [morally] wrong'. Ours however is effectively: 'a level of evidence sufficient for the belief to qualify for inclusion in the shared epistemic asset – ideally as knowledge'. This is neither circular nor trivial. By applying this criterion we discharge our duty to preserve, protect and develop our shared epistemic asset. It incorporates a moral demand to universalise what we recognise as our relevant epistemic context and our legitimate practical interest,¹⁴⁷ and thereby avoids the problematic agent-relativity implicit in individual practical interest.¹⁴⁸ And that same moral demand to universalise¹⁴⁹ drives why we would even care about the 'particular circumstances' and 'concerns' of our fellow 'inquirers'.¹⁵⁰

To the complaint that our criterion is too stringent to apply to all beliefs within the scope of EP4 we respond that as EP4 is a moral principle the criterion is ultimately aspirational. However that aspiration is in practice non-negotiable to the extent that the very possibility of knowledge and belief requires the existence of an epistemic community, which in turn requires our shared practical interest in preserving our shared epistemic asset.¹⁵¹ We could therefore see this as a minimum criterion for avoiding *insufficient* evidence: which would be to ensure we do not form beliefs in such a way as to undermine our epistemic community.¹⁵²

¹⁴⁶ p251.

¹⁴⁷ p249ff.

¹⁴⁸ p267.

¹⁴⁹ See p188.

¹⁵⁰ p265.

¹⁵¹ See p266.

¹⁵² The question might remain whether Craig's concept of a 'good informant ... whatever the particular circumstances of the inquirer' (1990, 1999:91) should be interpreted in a contextualist or invariantist way when understood as providing an account of *sufficient* evidence. Hawthorne's (2004:3) 'Syracuse' case (see p233) for example suggests we have shifting standards in the way we apply the concept of knowledge: would Craig's 'good informant' be happy to say he knows that (i) he will be in Syracuse next summer but not know that (ii) he will not die of a heart attack in the interim, even though (i) implies (ii)? But what is key for EP4 is that it provides a clear and unshifting standard of *insufficient* evidence given our common interest in protecting our shared epistemic asset.

Chapter 6 | In conclusion

6.2 Summary

This dissertation was prompted by the suspicion that, despite weaknesses, Clifford's principle (CP) had got something fundamentally right.¹⁵³ The aim was to tease out what that something might be.

CP is expressed in universal, absolutist terms:

...it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence (1877, 1879:186)

These expose it to attacks which a more restricted, qualified evidence principle might resist.

We gave two distinct but related reasons for excluding prescriptive beliefs. One was the challenge of explaining how evidence can support purely prescriptive beliefs.¹⁵⁴ The other was that applying the principle to prescriptive beliefs seemed incoherent – and not just because, if it was adopted it would be as a prescriptive belief, an evidence principle would have to presuppose itself in order to obey itself.¹⁵⁵

We then explored the idea of an evidence principle restricted to descriptive beliefs. A series of potential objections, clarifications and qualifications¹⁵⁶ culminated in:¹⁵⁷

EP4 {If anything is morally wrong, then} it is [morally] wrong [within the category of descriptive belief] to believe anything [knowingly or irresponsibly] on insufficient evidence {in the absence of any conflicting and overriding moral imperative} [except when the

¹⁵³ p40.

¹⁵⁴ p62ff.

¹⁵⁵ p72ff.

¹⁵⁶ p74ff.

¹⁵⁷ p92.

unjustified believing is outside the believer's voluntary control].

This makes explicit what CP left implicit: that it is a *morally* prescriptive principle. It therefore differs from a conceptual or 'intrinsic' ethics of belief like, say, Adler's (2002).¹⁵⁸ But even a conceptual requirement for sufficient evidence will generate expectations in a social, and therefore moral, context;¹⁵⁹ and Adler himself admits an 'ethical grounding' to the condition of 'full awareness' which underpins his account (Adler, 2002:9).¹⁶⁰

We then tested EP4 against James's very different focus on the conditions for legitimate over-belief in general and religious over-belief in particular.¹⁶¹

Contrasting moral positions are notoriously hard to prove or disprove. Complicating the case of James versus EP4 is a clash between two different existential perspectives and orientations. James prioritises individual liberty,¹⁶² whereas EP4, like CP, foregrounds the social and communal consequences of epistemic behaviour.¹⁶³ If EP4 does not conquer, it seems at least to survive unscathed.¹⁶⁴

The final step was to articulate how evidence can be sufficient to justify belief.¹⁶⁵ We summarised this at the end of the previous chapter¹⁶⁶ so only need repeat the eventual criterion:

Wherever possible, a level of evidence sufficient to support the level of justification required to be 'a good informant ... whatever the particular circumstances of the inquirer' (Craig, 1990, 1999:91)

¹⁵⁸ See p24ff.

¹⁵⁹ p35ff.

¹⁶⁰ p38ff.

¹⁶¹ p129ff, particularly p140ff.

¹⁶² Eg p181ff.

¹⁶³ p35ff and p210ff.

¹⁶⁴ p201ff.

¹⁶⁵ p205ff.

¹⁶⁶ p267ff.

6.2 Status of the principle

Our eventual criterion of evidential sufficiency is consistent with the social and moral orientation of EP4, and expands it into something which could be recommended to guide responsible epistemic behaviour. But recommended as what? What kind of thing is EP4? What is its status?

Its status is likely to be linked to the potential reasons we might have for holding and applying it. If for example EP4 was an empirical generalisation empirical evidence would count among its reasons for acceptance. If intended as an epistemological principle, perhaps a guideline in the methodology or good practice of belief, it might rest on a utilitarian cost-benefit analysis. If proposed as a synthetic a priori principle like Kant's Categorical Imperative we might expect a supporting deductive argument.¹⁶⁷

We have positioned EP4 throughout as a moral imperative like 'thou shalt not lie', so the reasons for holding it may exclude, or go beyond, empirical evidence. Indeed those reasons may be difficult or even impossible to state outside an overarching meta-ethical position.

If so we might not expect EP4 to appeal universally, but it could be seen as aligned with, say, utilitarianism, Kantianism or Golden Rule theory. EP4 presupposes not just morality itself, but also the individual's duty to preserve her moral sovereignty.¹⁶⁸

It does not completely outlaw religious belief. But it provides a framework for evaluating religious belief, so as to filter out negative and potentially dangerous consequences – for example when repressive prescriptive beliefs get bolted onto descriptive metaphysical over-beliefs.¹⁶⁹ It would therefore clash with divine command theory which, at least in principle, presupposes legitimate over-belief in the existence of a divine source of binding moral imperatives.

¹⁶⁷ But see **Appendix A5.1 Transcendental argument** below p292.

¹⁶⁸ p185ff.

¹⁶⁹ p136.

A positive conclusion is that we have not found a convincing reason for dismissing EP4. It is not incoherent in itself, nor does it conflict with other commonly accepted moral principles, for example about not lying, deceiving or hurting the innocent. (At the level of an individual action different principles considered in isolation may recommend conflicting decisions – for example it may be impossible in a particular context to tell the truth or obey EP4 without hurting someone’s feelings – but this applies to all moral principles.¹⁷⁰)

On the other hand we have not established any conclusive reason why you *must* hold EP4 – if for example you saw no value in shared social utility, and therefore no benefit in protecting and enhancing the social asset of shared knowledge. But again in this respect EP4 does not compare unfavourably with other familiar moral principles.

¹⁷⁰ p22.

Appendix A1 | Clifford on the evolution of conscience

In 'Right and Wrong: The Scientific Ground of Their Distinction' (1875, 1879)

Clifford asks:

'What is the best conscience? Or what ought I to think right?' (1875, 1879:170)

This question, he thinks,

resolve[s] itself into a question about the purpose or function of the conscience—why we have got it, and what it is good for (1875, 1879:166).

His functional approach licenses him to hop between a moral 'ought' and an 'is' of evolutionary explanation. Inspired by Darwin's *Descent of Man* (1871) he conjectures that human conscience

has been evolved and preserved because it is useful to the tribe or community in the struggle for existence against other tribes, and against the environment as a whole. The function of conscience is the preservation of the tribe as a tribe. And we shall rightly train our consciences if we learn to approve those actions which tend to the advantage of the community in the struggle for existence. (1875, 1879:167)

Following Aristotle's definition of an organism as 'that in which the part exists for the sake of the whole', where 'the shape and nature of the part are determined by the wants of the whole', he continues:

Society is [such] an organism, and man in society is part of an organism according to this definition, in so far as some portion of the nature of man is what it is for the sake of the whole—society. Now conscience is such a portion of the nature of man, and its function is the preservation of society in the struggle for existence. (1875, 1879:168-9)

Clifford takes natural selection to operate at the level of the individual and at the level of the 'tribes, clans, families, nations, towns' (1875, 1879:169) he belongs to. Accordingly,

for the purpose of the conscience the word community at any time will mean a group of that size and nature which is being selected or not selected for survival as a whole. Selection may be going on at the same time among many different kinds of groups. And ultimately the moral sense will be composed of various portions relating to various groups, the function or purpose of each portion being the advantage of that group to which it relates in the struggle for existence. (1875, 1879:169-70)

This only partially answers our question about the 'best conscience' – in the sense of what is most advantageous in the struggle for existence. We cannot even answer at the level of the individual organism 'because the organism grows in consequence of the struggle, and develops new wants while it is satisfying the old ones'. However 'right is an affair of the community, and must not be referred to anything else', because '[t]he first principle of natural ethics ... is the sole and supreme allegiance of conscience to the community' (1875, 1879:171-2):

there are no self-regarding virtues properly so called; those qualities which tend to the advantage and preservation of the individual being only morally right in so far as they make him a more useful citizen.

Clifford endorses utilitarianism to the extent that it 'explicitly sets forth the community as the object of moral allegiance' (1875, 1879:173) but rejects its classic definition of 'the end of right action' as the 'greatest happiness of the greatest number':

My happiness is of no use to the community except in so far as it makes me a more efficient citizen; that is to say, it is rightly desired as a means and not as an end. The end may be described as the greatest efficiency of all citizens as such. No doubt happiness will in the long run accrue to the

community as a consequence of right conduct; but the right is determined independently of the happiness...

Clifford's ultimate position seems to be therefore that the 'best conscience' ('moral sense') for an individual is what best achieves its 'purpose or function' – which is to promote survival at the level of his community. There appear to be four potential issues with this.

The first Clifford acknowledges: an individual may belong to 'many different kinds of groups'. His response seems to be that the broader the group, the higher the priority:

[P]art of the nature of a smaller group may be what it is for the sake of a larger group to which it belongs; and then we may speak of the function of the smaller group. ... [And] we may say that the function of the family is to promote the advantage of the nation or larger society... (1875, 1879:170)

The second is a potential implication of Clifford's position. We will take an evolutionary explanation of moral worth to identify what is right with what promotes the survival of the group the individual belongs to. By extension we can then identify what is right with what promotes the survival of the fittest members of the group at the expense of the least fit: other things being equal, this would promote the survival of the group in competition with other groups. But this could encourage measures significantly more brutal than Clifford's Poor Law policy.¹⁷¹ For example Herbert Spencer, whom Clifford admired,¹⁷² rubbished the idea of finding work for the unemployed: 'poor-law' theories asserting 'a man's right to a maintenance, and ... to have work provided for him' were 'erroneous' (1851:315). Better for nature to take its course:

[Predators] not only remove from herbivorous herds individuals past their prime, but also weed out the sickly, the malformed, and the least fleet or

¹⁷¹ See p55.

¹⁷² Eg: ...*Mr. Herbert Spencer, who has done so much for the whole doctrine of evolution and for all that is connected with it...* (Clifford, 1877, 1879:281)

powerful. By the aid of which purifying process ... the maintenance of a constitution completely adapted to surrounding conditions, and therefore most productive of happiness, is ensured. (1851:322)

This 'progress' holds throughout nature's 'higher creation', with its 'consummation' in humans, where the same

felicity-pursuing law ... never swerves for the avoidance of partial and temporary suffering. The poverty of the incapable, the distresses that come upon the imprudent, the starvation of the idle, and ... shouldering aside of the weak by the strong ... are the decrees of a large, far-seeing ... beneficence which brings to early graves the children of diseased parents, and singles out the low-spirited, the intemperate, and the debilitated as the victims of an epidemic. (1851:322-3)

The third issue arises from a more dispassionate analysis of natural selection: 'survival of the fittest' is profoundly tautologous, since evolutionary fitness is purely a measure of survival. What survives, and what promotes survival, depends on context. In some contexts increased fertility promotes survival; in others over-breeding leads to extinction. Some contexts favour increased complexity ('progress'); others favour regression to simpler, even degenerate, states. The only 'excellence' which favours evolutionary survival is excellence at survival.

Linking moral worth to survival is therefore problematic. If what is right for an individual is what promotes the survival of the group the individual belongs to, there can be no 'right' to arbitrate between two 'right' individuals from competing groups A and B. That is unless – and this seems to be Clifford's view – favouring what is 'right' for group A over what is 'right' for group B better promotes the survival of the group which groups A and B both belong to.

The fourth issue is the seemingly unwarranted assumptions linking 'is' to 'ought'. Spencer appears to treat evolutionary 'progress' as both a fact about the living world, if not the universe as a whole, and a moral target towards which conscious endeavour should be directed. Clifford does something similar, but with the

survival imperative. He seems to assume the sense of obligation we happen to experience as part of our evolutionary survival kit ('is') can explain and justify our actual obligation ('ought').¹⁷³ But this assumption is equally unwarranted, because we can always ask if we ought to feel the sense of obligation we happen to feel. With this assumption we could reduce morally prescriptive beliefs to hypothetical prudential beliefs, with hidden antecedents like 'if you want [the group you belong to] to win in the survival of the fittest, then...'. We have seen¹⁷⁴ that hypothetical prudential beliefs can be reworded as descriptive beliefs without significant loss of meaning. This could be why Clifford is happy to apply his principle to prescriptive beliefs – indeed to moral beliefs. However he seems to equate an evolutionary account of our sense of obligation ('is') with a justification of our actual obligation ('ought'), giving him an unwarranted licence to treat morally prescriptive beliefs as hypothetical prudential beliefs. But this route is not open to us, as the beliefs which worry us are categorical prescriptive beliefs, which by definition cannot be interpreted as hypothetical prescriptive beliefs with hidden antecedents.

¹⁷³ See p55.

¹⁷⁴ p44ff.

Appendix A2 | Bayes' Rule, evidence and prescriptive beliefs

Chapter 2¹⁷⁵ argued that a prescriptive belief may not be the sort of thing which could have supporting evidence. This was one of the reasons for excluding prescriptive beliefs from EP4. In this Appendix we will explore another possible reason for excluding prescriptive beliefs from EP4, based on how Bayes' Rule interprets the relationship between evidence and hypotheses, which we were not in a position to discuss in **Chapter 2**.

Consider this descriptive hypothesis:

H_d = 'The latest common ancestor of bonobos and humans was more recent than the latest common ancestor of gorillas and humans.'

Our evidence could be:

E_d = 'Humans and bonobos have 99% of DNA in common, whereas humans and gorillas have 96% of DNA in common.'

Chapter 2 used an 'initial working definition' of evidence:

... statement s counts as evidence for a belief that p if s is true and s provides objective reasons for thinking p is true or is likely to be true.¹⁷⁶

Chapter 5 then used Bayes' Rule to enrich the relationship between evidence and belief by introducing the probability of E_d if H_d were true = $\Pr(E_d/H_d)$.

Our background knowledge of evolution and genetics would lead us to assess $\Pr(E_d/H_d)$ as fairly high – certainly higher than if the percentages were reversed. More importantly it seems an expression we could evaluate, however imprecisely. If I was asked how likely it would be that humans and bonobos shared more DNA

¹⁷⁵ See p46ff.

¹⁷⁶ p62.

than humans and gorillas did, given that bonobos and humans shared a more recent common ancestor than gorillas and humans did, I might balk at an exact quantification. But I would understand the question, and be confident that E_d would be far more likely than if the percentages of shared DNA were the other way round.

Now consider a prescriptive hypothesis H_p , with proposed supporting evidence E_p :

H_p = 'Children should not be subjected to corporal punishment.'

E_p = 'Out of two equal populations of adults, one of which (sample A) had been subjected to corporal punishment as children while the other (sample B) had not, the incidence of individuals convicted of crimes of violence in sample A was three times that in sample B.'

Against our initial working definition E_p might seem to qualify as evidence for H_p . If I think H_p is something which could be true or false then the truth of E_p may be a reason for holding H_p as true (as the truth of E_d may be a reason for holding H_d as true).

But if we try to assess the degree to which E_p supports H_p by using Bayes' Rule we would need a value for $\Pr(E_p/H_p)$, the probability of E_p if H_p were true. But this seems obscure. If I was asked how likely I thought it was that the incidence of violent crimes committed by a sample of adults subjected to corporal punishment as children was three times that of an equivalent sample of adults not subjected to corporal punishment as children, given that children should not be subjected to corporal punishment, I am not sure I would understand the question. It would seem back to front, as if the questioner meant to ask how confident I would be in thinking that children should not be subjected to corporal punishment, given those relative rates of criminal activity, but muddled her words up.

Assuming we can talk of the 'truth' of a prescriptive hypothesis like H_p , its 'truth' does not give us a reason for expecting E_p to be true – in the way the truth of H_d does give us a reason for expecting E_d to be true. This is not to deny that a more

general *descriptive* hypothesis may be both a reason for holding prescriptive belief H_p and a reason for expecting E_p to be true, for example:

H_{gd} = 'Children subjected to corporal punishment are more likely to display aggressive behaviour in later life'.

So assuming the Bayesian conception of conditional probability is sound, this seems to give extra support to the idea that prescriptive beliefs are not the sort of thing that can be supported by evidence.

Compare now a (descriptive) utilitarian hypothesis:

H_u = 'Subjecting children to corporal punishment reduces rather than maximises net human happiness.'

The equivalent probability of E_p if H_u were true = $\Pr(E_p/H_u)$.

The expression $\Pr(E_p/H_u)$ seems as coherent as $\Pr(E_d/H_d)$, and we might assess the probability of E_p given the truth of H_u as quite high, especially if we add additional descriptive assumptions relating the incidence of violent crimes to net human happiness.

A utilitarian might see a factual claim like H_u as *justifying* a prescriptive claim like H_p . But a utilitarian claiming the 'truth' of H_p gives her a reason for expecting E_p to be true would have to see H_p and H_u as *identical* in propositional content – a rather more controversial claim.

For completeness consider a hypothetical prescription:

H_{hp} = 'If you want to avoid criminality, then children should not be subjected to corporal punishment'

If we apply the same evidence statement E_p then the probability of E_p if H_{hp} were true = $\Pr(E_p/H_{hp})$. This also seems coherent, unlike $\Pr(E_p/H_p)$. This is because the

overall content of H_{hp} is descriptive rather than prescriptive, making a descriptive claim about what achieves or causes what. It is like 'If you want your plants to live, you should water them', which has descriptive content equivalent to, say, 'Plants will live only if they are watered'.¹⁷⁷

A hypothetical prescriptive belief like 'If you want your plants to live, you should water them' is not the kind of prescriptive belief excluded from EP4, as clearly it can have supporting evidence. But a categorical prescriptive belief like H_p is excluded.

Assuming the Bayesian approach is sound, the asymmetry discussed here looks like another reason for thinking (categorical) prescriptive beliefs cannot have supporting evidence – at least in the straightforward way descriptive beliefs can.

¹⁷⁷ See also **2.2 Hypothetical prescriptive beliefs**, pp44–46.

Appendix A3 | EP4 and subjective Bayesian permissiveness

EP4 departs from standard Bayesian positions in two important respects. **Chapter 5** covered the prescribed currency of shared social utility, particularly in relation to shared duties towards our shared epistemic asset, rather than leaving ‘personal utility’ unqualified.¹⁷⁸ Another big difference is that a subjective Bayesian would (like James) allow far more permissiveness over beliefs and credences than our now contextualised EP4 would permit.

We will consider two scenarios. Both are where someone has a high credence for a hypothesis even though there appears no evidence (or insufficient evidence) for it. Both could be seen as legitimate from a subjective Bayesian perspective. However they fare differently against EP4.

- (A3.1) ‘Maverick’: The evidentially unsupported or under-supported hypothesis is logically connected to many other hypotheses but the individual has a consistent set of credences for the whole set. Some or all of these credences may well be very different from most other people’s.

Examples could include the belief that the world is flat,¹⁷⁹ or that the Apollo moon landings were a hoax.¹⁸⁰ In the Apollo case related hypotheses might include generally accepted ones about the contents and behaviour of the solar system; supporting ‘conspiracy’ hypotheses about NASA and the Nixon administration; and rejected ‘official’ explanations for apparent photographic anomalies like missing stars and the rippling flag.

- (A3.2) ‘Free floating’: The evidentially unsupported or under-supported hypothesis is logically connected with very few if any other

¹⁷⁸ See p249ff.

¹⁷⁹ See eg <http://www.theflatearthsociety.org/cms/index.php>. Accessed 12 May 2016.

¹⁸⁰ See eg https://en.wikipedia.org/wiki/Moon_landing_conspiracy_theories. Accessed 12 May 2016.

hypotheses. Again the set of credences is consistent, but this time it is because the set is very small and unrelated to most other hypotheses which the individual may have evidentially supported credences for.

An example could be James's 'religious hypothesis' (1896, 2000:214) which he sees as legitimate if it arises from a 'genuine option that cannot by its nature be decided on intellectual grounds' (1896, 2000:205).

The 'maverick' individual in (A3.1) may not necessarily be disobeying EP4. He can claim some evidence, and could have appropriately high credences for supporting hypotheses offering explanations for apparent counter-evidence suggesting that, say, the earth is spherical or that chemical analysis of Apollo moon rocks suggests a lunar source. To hold an unorthodox belief or set of beliefs is not to disobey EP4 – otherwise EP4 would stand in the way of much scientific progress. Sincere flat-earthers and Apollo hoax theorists could no doubt justifiably claim they are not knowingly or irresponsibly believing on insufficient evidence, which EP4 forbids. Bayesian agents may be assumed to be logically omniscient¹⁸¹ but EP4 does not impose it as a requirement.

But the maverick may not be so innocent. If he was genuinely aware of a balanced cross-section of available evidence and as a result knew or believed the evidence for his 'maverick' belief was insufficient, but insisted on holding the requisite credence values for a set of logically related hypotheses including the maverick hypothesis, then he would be disobeying EP4. Measured against EP4 he would be behaving irresponsibly with respect to belief (and as such endangering or devaluing the shared social asset of knowledge), even though a subjective Bayesian might say his high credence is legitimate because none of his related (possibly also 'maverick') credences constrained it to a lower value.

There is however no contradiction between being (i) licensed to have a certain credence from a subjective Bayesian point of view, where the avoided penalty for

¹⁸¹ See p219.

an incorrect (ie inconsistent) credence is to lose out in terms of personal utility specifically related to sure-loss contracts in respect of any imaginary gambles used to quantify credences; and being (ii) morally blameworthy for holding a certain belief or believing in a certain way (including having a particular credence for a hypothesis) where such behaviour could endanger or devalue the shared epistemic asset.¹⁸² You could be blameworthy under (ii) for availing yourself of the 'epistemic licence' granted under (i).

It may be that, in relation to the total set of evidence the agent chose to consider, his high credence for the maverick hypothesis was coherent against possible synchronic constraints set by logically related credences. If these credences are translated into imaginary gambles, the agent would not be exposing himself to sure losses. So even if the utility 'currency' were shared social utility, there would be no risk of a 'Dutch Book' loss. However deliberate 'partisan' selection of evidence to support chosen credences does not let the agent off the hook. In the broader context of normative decision theory in relation to epistemic behaviour, refusing to consider or look for any evidence is a special case of deliberate partisan selection of evidence¹⁸³ – which could expose the agent to the risk of losing bets which do not need to be lost and/or accepting bets at inappropriate odds. Such bets could risk loss of utility: which for EP4 we are quantifying in the currency of shared social value, and specifically that of a shared epistemic asset.

A similar argument applies in the case of (A3.2). From a subjective Bayesian perspective an agent may be licensed to hold, or have a high credence for, a specific 'free-floating' belief unconnected with any other credences which might otherwise constitute an inconsistent set and therefore expose the believer to sure-loss contracts in respect of imaginary gambles. But the agent may still be morally blameworthy by believing irresponsibly, and therefore not protecting the shared epistemic asset.

¹⁸² Even from a perspective of purely individual personal utility, there is no contradiction between (i) and (iii) losing out in terms of personal utility as a result of believing something which turns out to be false. The maverick could for example hold a credence of .95 for the hypothesis that on New Year's Day he will be able to fly, while ensuring that any other credences he holds for logically related hypotheses do not expose him to sure-loss contracts in respect of relevant imaginary gambles. On the first of January he launches himself off a roof and breaks his neck.

¹⁸³ See footnote 63 p229.

In both cases the agent could be accused of 'cheating': taking a free ride.¹⁸⁴ As a member of the epistemic community he would benefit from the shared epistemic asset, while asserting his freedom to endanger or undermine it by fostering credulity and gullibility as a result of believing on inadequate evidence.

Depending on the precise details of the story, the beliefs of Xhosa prophetess Nongqawuse which led to the 1856-7 mass destruction of crops and cattle¹⁸⁵ could qualify as either 'maverick' or 'free-floating'. Either way the shared epistemic asset of agricultural wisdom, from which Nongqawuse and those who spread her prophecies presumably benefited all their lives up to this point, was tragically compromised and overturned.


¹⁸⁴ See p254.

¹⁸⁵ See eg (Wikipedia, 2016). There are some parallels with Clifford's story about the 'medicine-man in Central Africa' (1877, 1879:200)

Appendix A4 | Bertrand's paradox

Bertrand's paradox (Bertrand, 1889) can be illustrated like this:

*Consider an equilateral triangle ABC inside a circle.
Draw a chord XY at random. What is the probability
that chord XY is longer than the side of triangle
ABC?*




The question seems to have (at least) three different answers, depending on how we randomly draw chord XY.

We could for example select random points on the circumference; or randomly centre the chord on a randomly selected diameter; or randomly select the mid-point of the chord.

A4.1 Three possible answers


A4.1.1 Random points on the circumference

Randomly select point X on the circumference to be one end of the chord. Draw equilateral triangle ABC inside the circle such that one of its vertices (say A) coincides with point X. The other end of the chord (Y) must be on one of three arcs AB, BC and CA. Because all three arcs are the same length it should be equally probable for Y to fall somewhere on arc AB or somewhere on arc BC or somewhere on arc CA. If, as in chord XY₁, Y falls on either arc AB or arc CA, the chord will be shorter than the side of the triangle. If Y falls on arc BC, as in chord XY₂, the chord will be longer than the side of the triangle. It is therefore twice as probable that the chord is shorter than that it is longer. Hence the probability that chord XY is longer than the side of triangle ABC is $\frac{1}{3}$.



A4.1.2 Randomly centred on randomly selected diameter


Randomly select a diameter AD. Draw equilateral triangle ABC inside the circle such that base BC is perpendicular to diameter AD. Then draw another circle inside triangle ABC. (This is the *incircle* of triangle ABC, whereas the original outer circle is the *circumcircle*.) Randomly select a point on diameter AD as the midpoint of chord XY also perpendicular to AD. If the chord overlaps the incircle (as in X_1Y_1) it will be longer than BC (the side of the triangle). If the chord does not overlap the incircle (as in X_2Y_2) it will be shorter than BC.



For any equilateral triangle the radius of its circumcircle is twice the radius of its incircle. So the probability that chord XY (eg X_1Y_1) overlaps the incircle (and is therefore longer than the side of the triangle) equals the probability that chord XY (eg X_2Y_2) does not overlap the incircle (and is therefore shorter than the side of the triangle). Hence the probability that chord XY is longer than the side of triangle ABC is $\frac{1}{2}$.

A4.1.3 Randomly selected midpoint

Consider the same equilateral triangle ABC with its circumcircle and incircle. Randomly select points within the circumcircle. Treat each point as the midpoint Z of a chord XY. A subset of these points will, like Z_1 , be within the incircle. They will be midpoints of chords like X_1Y_1 which are longer than the side of triangle ABC. The rest of the points will, like Z_2 , be outside the incircle, and be midpoints of chords like X_2Y_2 which are shorter than the side of triangle ABC.



The radius r_1 of the incircle of equilateral triangle ABC is half the radius r_2 of the circumcircle. Therefore the area of the incircle (πr_1^2) is a quarter the area of the circumcircle (πr_2^2). The probability of any midpoint Z being inside the incircle

rather than outside it is therefore $\frac{1}{4}$. Hence the probability that any chord XY is longer than the side of triangle ABC is $\frac{1}{4}$.

A4.2 Discussion

The three scenarios A4.1.1–3 above seem to show that the ‘randomly selected’ chord XY can be seen as a member of (and therefore randomly selected from) three different sets, based on the indefinite division of the circle’s circumference, diameter and area respectively. Each set yields a different probability for the chord being longer than the side of the triangle: $\frac{1}{3}$, $\frac{1}{2}$ and $\frac{1}{4}$ respectively. We cannot fix the probability in relation to the randomly selected chord without knowing which indefinite set it is randomly selected from, and therefore which randomising process is being applied.

Appendix A5 | Future directions

This dissertation could not explore every direction it wanted to, so below are suggestions for future work.

A5.1 Transcendental argument

Features of EP4 and its context suggest there might be a potential Kantian-style 'transcendental' argument for EP4 – although it may not be a purely deductive one. Lack of space restricts us to a mere sketch of what it could entail.

It would be an argument in support of EP4 itself, drawing on the justification for the proposed criterion of sufficient evidence,¹⁸⁶ and revolving around the shared epistemic asset of knowledge and justified belief. On our interpretation Craig's account of knowledge (and therefore, we argued, of justified belief¹⁸⁷) presupposes a context of community and intersubjectivity: interacting individuals with shared interests and values, acknowledging and assuming each other's subjectivity. We would need to argue that that context of community and intersubjectivity requires and rests on the community members' shared epistemic asset. Without a shared body of knowledge and justified belief the community members might struggle to interact sufficiently so as to acknowledge each other's subjectivities and understand and acknowledge each other's claims to knowledge and justified belief.

Assuming this move can be made, then for knowledge and justified belief to be possible – for 'all practical' purposes (Craig, 1990, 1999:98) at least – we must presuppose a shared epistemic asset. For that shared epistemic asset to survive, members of the community must, in turn, have a shared practical interest in its survival. Their awareness of shared practical interest would manifest as shared recognition of shared duties to ensure the survival of their shared epistemic asset.

¹⁸⁶ See p266.

¹⁸⁷ See p262.

On a day-to-day level these shared duties would include, for example, taking responsibility for the implications of what and how one believes, and whether one is doing so in 'full awareness' (Adler, 2002:9).¹⁸⁸ And as a measure of how shared their practical interests are they should share a broadly similar understanding of what 'all practical' purposes entail.¹⁸⁹

A5.2 Scientific revolutions

When distinguishing between descriptive and prescriptive beliefs there is a temptation to claim scientific belief as a paradigm case of descriptive belief, and scientific evidence as an exemplar of supporting evidence. Within the realm of purely descriptive beliefs we would want to compare a descriptive belief which seems to pass CP or EP4 by having adequate supporting evidence with one which fails to. We might therefore want to offer scientific belief as a standard of successful descriptive belief.

However the relationship between evidence and scientific belief may not prove unproblematic. For example during transitions (revolutions) in scientific belief there can be disagreement as to what counts as evidence, how evidence bears on statements of scientific theory, and how one should infer those statements from the evidence.¹⁹⁰ The disagreement could be seen as a clash between different methodological beliefs – therefore between different (non-moral) second-order prescriptive beliefs. If those prescriptive beliefs fall outside EP4, what justification would we have for choosing one rather than another?

The issue of underdetermination of theory by evidence merits much lengthier discussion than space here would allow. Future work would try to clarify exactly how and why a principle like EP4 expressed in moral language might be exposed by considerations of underdetermination and how it might be defended.

¹⁸⁸ See p35ff.

¹⁸⁹ See p262.

¹⁹⁰ Mentioned p63.

A5.3 Moral revolutions

'Moral revolutions' are another potential loose end, illustrated by the following thought experiment. A relatively happy but closed community are united in their shared over-beliefs about the nature and wishes of a particular pantheon: the Norse Æsir, say. It is a context where 'the only reason for belief is that everybody has believed the thing for so long that it must be true' (Clifford, 1877, 1879:200). Continued social cohesion may however depend on continued shared over-belief. This is not to rule out the possibility that even better times may follow from eventual loss of belief, but only after an interval of social disintegration and suffering.

The 'closed community' has implications for the interpretation of 'indefinite intersubjectivity'¹⁹¹ and 'indefinitely good informant'¹⁹² in **Chapter 5**. The individual members of the community may see a 'good informant' as one who describes natural phenomena in terms of the actions of Odin, Thor and Freyja. But 'indefinite intersubjectivity' implies projecting outside the community – into possibly alien circumstances and concerns which community members might never envisage.

Two immediate responses come to mind. The simpler is that this scenario may not challenge EP4, because of its 'conflicting and overriding moral imperative' caveat. The overriding moral imperative may not be to preserve the existing status quo come what may. But we might entertain an imperative to limit collateral damage by adopting careful, creative and gradualist ways to promote transition to more evidence-based belief.

A more complex response might explore whether a principle like EP4 relies on a particular social context. It could be a luxury which only 'enlightened' liberal democracies can enjoy or even contemplate.

¹⁹¹ p263.

¹⁹² p265.

A5.4 Moral encroachment

There has been much recent debate on 'moral encroachment', defined by Rimu (2019b:10) as the thesis that 'the epistemic justification of our beliefs can be determined, in part, by the moral demands of our situation, i.e., the moral stakes'. Example literature includes Basu (2019a), Basu (2019b), Bollinger (2018), Gardiner (2018), Moss (2018) and Schroeder (2018), many of which employ test cases involving possible racial profiling. Basu (2019a) gives the example of a 'supposedly rational racist' waiter called Spencer who is 'tired of constantly being called a racist' whenever he believes of a particular black diner that he will tip worse than the average white diner, based on his knowledge of 'studies that show that on average black diners tip substantially less than white diners.'

Because moral encroachment can be considered at least analogous to pragmatic encroachment, which we have discussed in the articulation of EP4 as a moral principle, it would be valuable to check if anything in the moral encroachment literature might lead us to revise our treatment of EP4.

At first sight though our transformation of the 'practical interest' of pragmatic encroachment into the 'practical interest we (morally) ought to have' means our account should already be sufficiently resourced to ensure that beliefs which are cases of high 'moral stakes' call for an appropriately high level of evidential support. So in the particular context of potential racial profiling our incorporation of Craig's 'good informant as to whether p whatever the particular circumstances of the inquirer' (1990, 1999:91) should accommodate inquirers who may for example correctly insist that the informant fully appreciates both the range of reference classes a potential target of racial profiling may belong to and the 'harms and wrongs that accompany being racially profiled' (Basu, 2019b). Spencer does not after all know what tip he will get until the diner pays the bill, and he will probably never know why the diner tipped the way he did.

Bibliography

- Adler, J.E. 2002. *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Adler, J.E. 2017. 'Epistemological Problems of Testimony', *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), Edward N. Zalta, Ed. Available: <https://plato.stanford.edu/archives/win2017/entries/testimony-episprob/> [23 July 2020].
- Aikin, S.F. 2014. *Evidentialism and the Will to Believe*. London and New York: Bloomsbury Academic.
- Alston, W.P. 1988. The deontological conception of epistemic justification. *Philosophical Perspectives*. 2:257-299.
- Audi, R. 1997. 'The Place of Testimony in the Fabric of Justification and Knowledge'. *American Philosophical Quarterly*, 34: 405-422.
- Ayer, A.J. 1956. *The Problem of Knowledge*. Harmondsworth: Penguin.
- Basu, R. 2019a. 'The Wrongs of Racist Beliefs'. *Philosophical Studies*, 176 (9): 2497-2515
- Basu, R. 2019b. 'Radical Moral Encroachment'. *Philosophical Issues*, 29:9–23.
- Belzer, M. 2007. 'Mere Stranger'. In *Philosophers without Gods: Meditations on Atheism and the Secular Life*. L.M. Antony, Ed. New York: Oxford University Press. 90-102.
- Bird, A. 2007. 'Justified judging'. *Philosophy and Phenomenological Research*. 74(1):81-110.
- Blackburn, S. 2007. 'Religion and Respect'. In *Philosophers without Gods: Meditations on Atheism and the Secular Life*. L.M. Antony, Ed. New York: Oxford University Press. 179-192.
- Bolinger, R. J. 2018. 'The rational impermissibility of accepting (some) racial generalizations'. *Synthese*. DOI: <https://doi.org/10.1007/s11229-018-1809-5>.
- Burge, T. 1993. 'Content Preservation'. *The Philosophical Review*, 102: 457–488.
- Burge, T. 1997. 'Interlocution, Perception, and Memory'. *Philosophical Studies*, 86: 21–47.

- Burnyeat, M.F. 1983. 'Can the Skeptic Live His Skepticism?'. In *The Skeptical Tradition*. M.F. Burnyeat, Ed. Berkeley: University of California Press. 117-148.
- Casscells, W., Schoenberger, A. & Graboys, T.B. 1978. 'Interpretation by Physicians of Clinical Laboratory Results'. *New England Journal of Medicine*. (299):999-1001.
- Chignell, A. 2018. 'The Ethics of Belief', *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.). Available: <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/> [3 July 2020].
- Chisholm, R.M. 1991. 'Firth and the Ethics of Belief'. *Philosophy and Phenomenological Research*. 51(1):119-128.
- Christensen, D. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford.
- Clarke, R. 2013. 'Belief is Credence One (In Context)'. *Philosophers' Imprint*. 13:1-18.
- Clifford, W.K. 1875, 1879. 'Right and Wrong: The Scientific Ground of Their Distinction'. In *Lectures and Essays, Volume II*. L. Stephen & F. Pollock, Eds. London: MacMillan and Co. 124-176.
- Clifford, W.K. 1877, 1879. 'The Ethics of Belief'. In *Lectures and Essays, Volume II*. L. Stephen & F. Pollock, Eds. London: MacMillan and Co. 177-212.
- Coady, C.A.J. 1992. *Testimony: A Philosophical Study*. Oxford: Oxford University Press.
- Craig, E. 1990, 1999. *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- Davidson, D. 1982. 'Rational Animals'. *Dialectica*. 36(4):317-327.
- Davidson, D. 2001a. 'The Emergence of Thought'. In *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press. 123-134.
- Davidson, D. 2001b. 'Three Varieties of Knowledge'. In *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press. 205-220.
- Diderot, D. 1875-1877. *Pensées philosophiques, LIX, Oeuvres*. Paris: .
- Douven, I. 2003. 'Nelkin on the Lottery Paradox'. *The Philosophical Review*. 112:395-404.
- Feldman, F. 2005. 'The Open Question Argument: What it Isn't; and What it Is'. *Philosophical Issues*. 15(1):22-43.

- Feldman, R. 2001. 'Voluntary Belief and Epistemic Evaluation'. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. M. Steup, Ed. New York: Oxford University Press. 77-92.
- Foley, R. 1993. *Working Without a Net: A Study of Egocentric Epistemology*. Oxford: Oxford.
- Frankish, K. 2009. 'Partial Belief and Flat-Out Belief'. In *Degrees of Belief (Synthese Library 342)*. F. Huber & C. Schmidt-Petri, Eds. UK: Springer. 73-93.
- Franklin, J. 1998. 'Two Caricatures, I: Pascal's Wager'. *International Journal for Philosophy of Religion*. (44):109-114.
- Fricker, E. 1987. 'The Epistemology of Testimony'. *Proceedings of the Aristotelian Society Supplementary*, 61: 57–83.
- Fricker, E. 1994. 'Against Gullibility'. In *Knowing from Words*. B.K. Matilal & A. Chakrabarti, Eds. Boston: Kluwer. 125-161.
- Fricker, E. 1995. 'Critical Notice: Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony'. *Mind*, 104: 393–411.
- Fricker, E. 2004. 'Testimony: Knowing through Being Told'. In *Handbook of Epistemology*. I. Niiniluoto, M. Sintonen & J. Wolenski, Eds. Dordrecht: Kluwer Academic Publishers. 109-130.
- Fricker, E. 2006. 'Testimony and Epistemic Autonomy'. In *The Epistemology of Testimony*. J. Lackey & E. Sosa, Eds. Oxford: Oxford University Press. 225-251.
- Gale, R.M. 1993. *On the Nature and Existence of God*. Cambridge: Cambridge University Press.
- Gale, R.M. 1999, 2007. *The Divided Self of William James*. Cambridge: Cambridge University Press.
- Gardiner, G. 2018. 'Evidentialism and Moral Encroachment'. In *Believing in Accordance with the Evidence: New Essays on Evidentialism*. K. McCain, Ed. Springer. 169-195.
- Haack, S. 1997, 2001. "The Ethics of Belief" Reconsidered. In *Knowledge, Truth, and Duty*. M. Steup, Ed. New York: Oxford. 21-33.
- Hacking, I. 1972. 'The Logic of Pascal's Wager'. *American Philosophical Quarterly*. 9(2):186-192.
- Hacking, I. 2001, 2009. *An Introduction to Probability and Inductive Logic*. New York: Cambridge University Press.

- Hall, R. 2011. 'The Clifford/James Debate'. *Analytic Teaching and Philosophical Praxis*. 31(1):79-88.
- Hansson, S.O. 2005. *Decision Theory: A Brief Introduction*. Available: <http://people.kth.se/~soh/decisiontheory.pdf> [24 August 2014].
- Harman, G. 1986. *Change in View*. Cambridge: MIT Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hazlett, A. 2010. 'The Myth of Factive Verbs'. *Philosophy and Phenomenological Research*. 80(3):497-522.
- Howard, R. 1966. 'Information Value Theory'. *IEEE Transactions on Systems Science and Cybernetics*. Vol. 2(1):22-26.
- Howson, C. & Urbach, P. 2006. *Scientific Reasoning: The Bayesian Approach*. 3rd ed. Chicago and La Salle, Illinois: Open Court.
- Hume, D. 1777, 1902, 1970. An Enquiry Concerning the Principles of Morals. In *Enquiries Concerning the Human Understanding and Concerning the Principles of Morals*. L.A. Selby-Bigge, Ed. Second ed. London: Oxford University Press. 169-323.
- Hume, D. 1969. *A Treatise of Human Nature*. Harmondsworth: Penguin.
- Hume, D. 1977 [1748]. *An Enquiry Concerning Human Understanding*. Eric Steinberg, Ed. Indianapolis: Hackett Publishing Company.
- Hume, D. 1978 [1740]. *A Treatise of Human Nature*. L.A. Selby-Bigge, Ed. Second edition, revised P. H. Nidditch (ed.), Oxford: Oxford University Press.
- Ichikawa, J.J. & Steup, M. 2001, 2017. 'The Analysis of Knowledge', *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition). Available: <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/> [21 January 2019].
- James, W. 1891, 2000. 'The Moral Philosopher and the Moral Life'. In *Pragmatism and Other Writings*. G. Gunn, Ed. London: Penguin. 242-263.
- James, W. 1896, 2000. 'The Will to Believe'. In *Pragmatism and Other Writings*. G. Gunn, Ed. London: Penguin. 198-218.
- James, W. 1897, 1912a. 'Is Life Worth Living?'. In *The Will to Believe: and Other Essays in Popular Philosophy*. New York: Longmans, Green, and Co. 32-62.
- James, W. 1897, 1912b. 'The Sentiment of Rationality'. In *The Will to Believe: and Other Essays in Popular Philosophy*. New York: Longmans, Green, and Co. 63-110.

- James, W. 1897, 2000. 'Preface' to *The Will to Believe and Other Essays in Popular Philosophy*. In *Pragmatism and Other Writings*. G. Gunn, Ed. London: Penguin. 193-197.
- James, W. 1902, 1960. *The Varieties of Religious Experience*. London: The Fontana Library.
- James, W. 1907, 2000. 'Pragmatism'. In *Pragmatism and Other Writings*. G. Gunn, Ed. London: Penguin. 5-132.
- Jeffrey, R. 1983. *The Logic of Decision*. 2nd ed. Chicago: University of Chicago Press.
- Kahneman, D. 2012. *Thinking, Fast and Slow*. London: Penguin.
- Kant, I. 1786. 'What does it mean to orient oneself in thinking?'. In *Immanuel Kant: Religion and Rational Theology*. A.W. Wood & G. Di Giovanni, Eds. Cambridge: Cambridge University Press. 1-18.
- Kant, Immanuel 2005. *The Moral Law: Groundwork of the Metaphysics of Morals*. London and New York: Routledge.
- Keynes, J.M. 1921. *A Treatise on Probability*. London: Macmillan.
- Kierkegaard, S. 1974. *Concluding Unscientific Postscript*. Princeton University Press.
- Kierkegaard, S. 2000. *The Essential Kierkegaard*. Princeton: Princeton University Press.
- Kornblith, H. 1986. 'Naturalizing Rationality'. In *Naturalism and Rationality*. N. Garver & P.H. Hare, Eds. Buffalo, New York: Prometheus Books. 115-133.
- Kyburg, H. 1961. *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Lecky, W. 1880. *History of European Morals from Augustus to Charlemagne*. 4th ed. London: Longmans, Green, and Co.
- Lewis, D. 1999. 'Why conditionalize?' In *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press. 403-407.
- Madigan, T.J. 2009. *W. K. Clifford and "The Ethics of Belief"*. Newcastle: Cambridge Scholars Publishing.
- Maher, P. 1993. *Betting on Theories*. Cambridge: Cambridge.

- Marshall, C., Lightman, B. & England, R., Eds. 2015. *The Papers of the Metaphysical Society, 1869-1880: A Critical Edition*. Oxford: Oxford University Press.
- McHugh, C. 2011. 'What do we aim at when we believe?' *Dialectica*. 65(3):369-392.
- McHugh, C. 2014. 'Exercising Doxastic Freedom'. *Philosophy and Phenomenological Research*. 88(1):1-37.
- Meiland, J.W. 1980. 'What Ought We to Believe? or The Ethics of Belief Revisited'. *American Philosophical Quarterly*. 17(1):15-24.
- Molière, J. 2003, 2012. *The Imaginary Invalid*. Project Gutenberg.
- Moore, G.E. 1903, 1922. *Principia Ethica*. London: Cambridge University Press.
- Moss, S. 2018. 'Moral Encroachment'. *Proceedings of the Aristotelian Society*. 118, 177-205.
- Nelkin, D.K. 2000. 'The Lottery Paradox, Knowledge, and Rationality'. *The Philosophical Review*. 109:373-409.
- Pascal, B. 1966. *Pensées*. London: Penguin.
- Peirce, C.S. 1878, 1998. 'How to Make Our Ideas Clear'. In *Charles S. Peirce: The Essential Writings*. E.C. Moore, Ed. New York: Prometheus Books. 137-157.
- Peirce, C.S. 1955. 'The General Theory of Probable Inference'. In *Philosophical Writings of Peirce*. J. Buchler, Ed. New York: Dover. 190-217.
- Plantinga, A. 1981. 'Is Belief in God Properly Basic?'. *Noûs*. 15(1):41-51.
- Plantinga, A. 1983. 'Reason and Belief in God'. In *Faith and Rationality: Reason and Belief in God*. A. Plantinga & N. Wolterstorff, Eds. Notre Dame, Indiana: University of Notre Dame Press. 16-93.
- Plantinga, A. 1993. *Warrant: The Current Debate*. New York: Oxford University Press.
- Plantinga, A. 2000. *Warranted Christian Belief*. New York: Oxford University Press.
- Plato 1955, 1974 (Revised). *The Republic*. 2nd ed. Harmondsworth, UK: Penguin.
- Quine, W. 1980. 'Two Dogmas of Empiricism'. In *From a Logical Point of View*. 2nd ed. Cambridge: Harvard University Press. 20-46.

- Ramsey, F. 1931. 'Truth and Probability'. In *Reprinted in Philosophical Papers*. D.H. Mellor, Ed. Cambridge: Cambridge University Press.
- Reid, T. 1983. *Inquiry and Essays*. R. Beanblossom & K. Lehrer. Eds. Indianapolis: Hackett.
- Rosen, G. undated. *Pascal's Wager*. Available: <http://www.princeton.edu/~grosen/puc/phi203/Pascal.html> [9 August 2014].
- Sartre, J. 2007. *Existentialism is a Humanism*. New Haven & London: Yale University Press.
- Schroeder, M. 2018. 'When Beliefs Wrong'. *Philosophical Topics*. 46(1):115-128.
- Searle, J.R. 1964. 'How to Derive "Ought" From "Is"'. *The Philosophical Review*. 73(1):43-58.
- Singer, P. 2005. 'Peter Singer "Tells It Like It Is"'. *Freethought Today*. Vol. 22, No. 4 (May 2005). Freedom From Religion Foundation, Inc. Available: <https://ffrf.org/publications/freethought-today/item/13232-peter-singer-tells-it-like-it-is> [24 July 2020]
- Singer, P. 2011. *The Expanding Circle*. Princeton NJ: Princeton University Press.
- Spencer, H. 1851. *Social Statics*. London: John Chapman.
- Stanley, J. 2005. *Knowledge and Practical Interests*. Kindle ed. Oxford: Clarendon Press.
- Strevens, M. 1998. 'Inferring Probabilities From Symmetries'. *Noûs*. 32(2):231-246.
- Strevens, M. 2005. 'The Bayesian Approach to the Philosophy of Science'. In *Macmillan Encyclopaedia of Philosophy*. Second ed. Macmillan.
- Talbott, W. 2013. 'Bayesian Epistemology', *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.). Available: <http://plato.stanford.edu/archives/fall2013/entries/epistemology-bayesian/> [13 March 2014].
- Taylor, K.A. 2007. 'Without the Net of Providence: Atheism and the Human Adventure'. In *Philosophers without Gods: Meditations on Atheism and the Secular Life*. L.M. Antony, Ed. New York: Oxford University Press. 150-163.
- Teller, P. 1973. 'Conditionalization and Observation'. *Synthese*. 26(2):218-258.
- Thorsrud, H. undated. *Internet Encyclopedia of Philosophy: Ancient Greek Skepticism*. Available: <http://www.iep.utm.edu/skepanci/> [24 May 2015].

- Tversky, A. & Kahneman, D. 1974. 'Judgment under Uncertainty: Heuristics and Biases'. *Science*. 185:1124-1131.
- Van Fraassen, B. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.
- Vogel, J. 1990. 'Are there Counterexamples to the Closure Principle?'. In *Doubting: Contemporary Perspectives on Skepticism*. M. Roth & G. Ross, Eds. Dordrecht: Kluwer.
- Wikipedia 2016. *Nongqawuse*. Available:
<https://en.wikipedia.org/wiki/Nongqawuse> [2 May 2016].
- Williams, B. 1973. 'Deciding to believe'. In *Problems of the Self*. Cambridge: Cambridge University Press. 401-419.
- Williamson, T. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.
- Winters, B. 1979. Believing at will. *The Journal of Philosophy*. 76(5):243-256.