


WHAT JUSTIFIES BELIEF?
PROBABILITY, NORMALCY, AND THE FUNCTIONAL THEORY

Marvin Backes

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews



2019

Full metadata for this item is available in
St Andrews Research Repository
at:
<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:
<http://hdl.handle.net/10023/17919>

This item is protected by original copyright

What Justifies Belief?
Probability, Normalcy, and the Functional Theory

Marvin Backes



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of
Doctor of Philosophy (PhD)
at the University of St Andrews

October 2018

Candidate's declaration

I, Marvin Backes, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 74,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree.

I was admitted as a research student at the University of St Andrews in September 2015.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date

Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date

Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Marvin Backes, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis: No embargo on print copy. No embargo on electronic copy.

Date

Signature of candidate

Date

Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Marvin Backes, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date

Signature of candidate

For my family, with love

Acknowledgments

Over the last three years in which this dissertation was written I have been fortunate to benefit from the help, guidance, and support of a large number of people. Sometimes this help took the form of valuable feedback and comments on my work. Other times it took the form of interesting discussions that ultimately helped me develop some of the ideas presented here. And yet other times, especially in moments of frustration and self-doubt - which any graduate student is likely to be familiar with - this help came in the form of reassurance and kind words of encouragement. As such, many people have, in some way or other, played a part in helping me with his project. And for that I thank you all!

Special thanks go to my two supervisors Patrick Greenough and Crispin Wright. Patrick has not just played an important role in shaping, sharpening, and improving the ideas presented here, but he also, more than anyone else, helped me develop my voice as a writer – one of the less talked about skills in philosophy. Sometimes this took the form of hour-long supervisions, and it was here that I learned some of the most valuable lessons of my time as a student. At the same time, he always encouraged me to think broadly and not to stop probing too soon. However, one of Patrick's greatest virtues as a supervisor, aside from being an excellent teacher and philosopher, is a genuine care and concern for his students that goes beyond their philosophical work. Thank you, Patrick.

I also owe special gratitude to Crispin Wright for providing great comments and helpful suggestions on my work. Few things are more conducive towards understanding more clearly one's own position and its underlying intricacies than conversations with Crispin. As will become apparent, his influence can be found in many places within this work.

I am also thankful to Matt McGrath, Sònia Roca-Royes, and Martin Smith for helpful discussions which have helped me to see various issues more clearly. Likewise, I'm indebted to Michael Popich, an inspirational teacher and the person

who first sparked my interest in philosophy. I also want to thank the AD Links Foundation who has provided funding for my studies over the last three years.

I am grateful to my friends and colleagues at St Andrews. Spending three years in this small town together, they have not just had a considerable influence on my philosophical thinking and on the ideas presented here, but they also made the past few years (especially Friday nights) very enjoyable. Thank you for that. Another special thanks goes to Taste in St Andrews – the café where much of this work was written – and the wonderful people whom I've met here over the years.

Finally, I want to thank my family – especially my mother Claudia – for their love, support, and encouragement over the years. Thank you.

Abstract

‘What justifies belief?’ This question is arguably one of the most important questions in contemporary epistemology. The first part of this study looks at two very different answers to the above question, but ultimately finds both of them wanting. According to *probabilistic* accounts of justification, the property that makes a belief justified is some property along the lines of *being highly probable*. I call this picture of justification the *Lockean View*. In contrast, according to the most prominent *non-probabilistic* accounts of justification, the property that justifies belief is some property along the lines of *being true in all normal worlds*. I call this non-probabilistic picture of justification the *Normalcy View*. However, as we will see, both families of views turn out to be problematic. While probabilistic accounts are incompatible with an attractive principle called *multi premise closure* (MPC), non-probabilistic accounts, I argue, are too demanding and therefore too stingy. This leaves us in a dilemma; neither probabilistic nor non-probabilistic accounts of justification seem to be wholly satisfactory. I call this the *(MPC)-Stinginess Dilemma*.

The second part of this study is concerned with how we should respond to this dilemma. After considering but rejecting some initial options, I argue that the dilemma can be avoided if we reject the almost universally accepted *monist* assumption that there is only one way for a belief to be justified; or, that there is only one property that can make a belief justified. Subsequently I develop and defend a novel, *pluralist*, theory of epistemic justification. I call it the *Functional Theory of Justification*. One upshot of the functional theory is that it makes room for the idea that there is more than one way for a belief to be justified; or, more precisely, that depending on our *epistemic environment*, justification can be realized by different properties.

Contents

Introduction	1
1. Epistemic Justification: A Minimalist Conception	15
2. The Orthodox Picture: The Lockean View and the Challenge from Statistical Evidence	41
3. Multi Premise Closure and the Paradox of the Bitter Pill	69
4. A Non-Probabilistic Alternative: Normalcy Views and the Problem of Easy-Defeat	97
5. Exploring New Territory: Pluralism about Justification and The Functional Theory	133
6. Applications: New Solutions to Our Epistemic Puzzles	167
Conclusion	197
Bibliography	201

Introduction

1. *Theories of Justification: A New Starting Point for a Taxonomy*

We all hold a large number of beliefs – e.g. that $2+2=4$, that our favourite football team is going to win the cup this year, that the bank will be open on Saturday, that we are not going to win the lottery, that we have hands, etc. But which of our beliefs are we actually justified in holding? In other words, which of our beliefs are, epistemically speaking, in good standing and which ones are not? Theories of *propositional justification*¹, which are the primary concern of this study, set out to answer precisely this question; more generally, their chief concern is to provide an answer to the question “what justifies belief?”

When looking at some of the prominent answers to the above question, one quickly gets the impression that there are broadly speaking two fundamentally opposing families of views – *internalist* and *externalist* theories of justification.

Internalists, as is widely known, hold that whether a belief is justified depends *entirely* on factors internal or accessible to the agent - usually the total body of evidence *available* to the agent. As such, internalists tend to approach the question of whether someone is justified in believing some proposition P from a *first-person perspective*. In contrast, externalist, as is also well known, reject internalism’s central claim that justification is an *entirely* internal matter. Instead, externalist accounts of justification take a *third-person perspective* and allow for factors external to the agent - and of which the agent may be entirely *unaware* - to enter the picture. The *locus*

¹ *Propositional justification* is often contrasted with *doxastic justification*. Propositional justification, as the name suggests, is concerned with the justificatory status of propositions – i.e. it is concerned with whether some agent S *is justified in*, or *has justification for*, believing some proposition P. Doxastic justification on the other hand is concerned with whether S *justifiably believes* that P. The two notions are usually taken to be closely related. According to the orthodox picture, broadly speaking, S’s belief that P is doxastically justified if (i) P is propositionally justified for S and (ii) S’s belief that P is *based* on the reason that make P propositionally justified for S (See, for instance, Feldman (2002) and Kvanvig (2003)). For a good discussion and criticism of the orthodox picture, see Turri (2010).

classicus for externalist theories of justification is Goldman's (1979) reliabilism, according to which a belief is justified only if it was produced by a historically reliable belief-forming method.²

While internalism was the default position for a long time the arrival of externalist views of justification effectively began to challenge internalism's default status by providing a seemingly compelling alternative that promised to solve a number of problems faced by their internalist counterparts.³ With two fundamentally opposing views on the stage, a lot of subsequent energy was spent on fleshing out the respective views in more detail and on identifying the problems and costs of the two approaches.⁴ Hence, the internalism/externalism divide looks like a natural starting point for taxonomizing theories of justification.

However, only because the internalism/externalism divide appears to be an *obvious* starting point for taxonomizing theories of justification, it is by no means the only available option. An equally plausible, and perhaps more fundamental, way of carving up justification's theoretical landscape is to distinguish between *probabilistic* and *non-probabilistic* theories of justification. By 'probabilistic' I do not just mean theories that are formulated using probabilistic language, but instead I use the term more broadly to pick out theories that explain justification in terms of *high probability*.⁵ On probabilistic theories, there usually exists some epistemic good-making feature, e.g. an agent's total body of evidence or the reliability of a belief-forming process, and a belief in P is justified as long as P's degree of probability,

² Whether or not a belief-forming method has been historically reliable is of course something external to the agent, and, importantly something that the agent has no epistemic access to.

³ See Bonjour (2002:235)

⁴ For a good overview of the debate see Vahid (2010), Madison (2010).

⁵ According to my proposed use of the term 'probabilistic' it is possible for a theory to be probabilistic even though it does not explicitly contain any probabilistic language. For instance, the view that a belief that P is justified iff P is true in most possible worlds would be a probabilistic one. Likewise, the view that a belief that P is justified iff the evidence sufficiently favours P would count as probabilistic. Both accounts can reasonably be understood as explaining justification in terms of *sufficient likelihood* or *high probability*.

Furthermore, according to my use of the term 'probabilistic' a theory can fail to be probabilistic despite the fact that it contains probabilistic language. For example, the view that a belief is justified iff its degree of probability is 1 would not be probabilistic in the relevant sense – it does not explain justification in terms of high probability.

given the relevant good-making feature, is above some threshold value t (where $t < 1$) required for justification. We may call this the Lockean View of justification.

Lockean View A belief that P is justified for S iff P is highly probable for S.

More specific versions of the Lockean View have been widely embraced by a large number of epistemologists;⁶ in fact, it seems appropriate to consider this probabilistic approach to justification as one of the leading paradigms of contemporary epistemology. Importantly, on the Lockean picture, the property that makes a belief justified is some property along the lines of *being highly probable*.

Non-probabilistic theories of justification, in contrast, are relatively new and have only recently started to receive more serious attention. Non-probabilistic theories, broadly construed, reject the claim that high probability is either *necessary* or *sufficient* for justification and instead explain justification by appeal to some other property. A key task of any non-probabilistic theory is of course to explain what this other property may be in virtue of which beliefs come to be justified. The most promising proposals have been developed and defended in recent work by Martin Smith (2010, 2016) and Jarrett Leplin (2009). On their views, broadly speaking, a belief that P is justified for S only if, given S's evidence, there exists no normal world in which S falsely believes that P. We may call this the *Normalcy View*.

Normalcy View The belief that P is justified for S only if, given S's evidence E, there exists no *normal world* in which S falsely believes that P.

Normalcy views are non-probabilistic because they do not explain justification in terms of high probability. Instead, according to the Normalcy View, the property that makes a belief justified is the property of *being true in all normal worlds*; or, for short, of *being normally true*.

⁶ For examples, see Goldman (1979), Bonjour (1985: 6), Alston (1988: 269), Moser (1989:42), Foley (1992), Plantinga (1993: 18), Swinburne (2001), Conee & Feldman (2004: 100), Pryor (2004: 352), Bergmann (2006: ch.6), and Sturgeon (2008).

This divide between probabilistic and non-probabilistic theories of justification provides an alternative starting point for taxonomizing theories of justification. And one might think that the probabilistic/non-probabilistic divide is in some way more fundamental than the internalism/externalism divide. Here is why: While the internalism/externalism divide focuses on the *location* of the epistemic good-making feature relative to the agent – i.e. whether it is internal or non-internal to the agent - the probabilistic/non-probabilistic divide focuses directly on the *kind* of property used to explicate justification. In other words, while internalists and externalists are in agreement that the property that makes a belief justified is something along the lines of *being highly probable* and disagree only about what we should look at in determining the relevant degrees of probability, the disagreement between proponents of probabilistic and non-probabilistic theories of justification is about the very *kind* of property we should use to analyse the notion of justification. Now, it seems plausible that the question of *what property to look for* is more fundamental than the question of *where to look for it*. After all we must first decide *what* we are looking for before we can decide on *where* to look for it. Now, I am not claiming that these considerations are conclusive, but I do think that they provide some reason for thinking that the divide between probabilistic and non-probabilistic theories of justification is more fundamental than the one between internalist and externalist theories. If this is correct, then the probabilistic/non-probabilistic divide provides an interesting new starting point for theorizing about the notion of justification. But even if one disagrees with the claim that the probabilistic/non-probabilistic divide is more fundamental than the internalism/externalism divide, or one thinks that these considerations of comparative fundamentality are altogether irrelevant, one can still appreciate the claim that for the purpose of taxonomizing theories of justification the probabilistic/non-probabilistic divide provides an interesting new starting point that is worth exploring.

2. *The Probabilistic/Non-Probabilistic Divide and a Dilemma*

In light of the extensive work on the internalism/externalism debate over the last decades epistemologists have arrived at a pretty thorough understanding of the

advantages, problems, and costs associated with the two approaches. For instance, internalism is widely thought to struggle with cases of forgotten evidence⁷, is sometimes accused of providing an over-intellectualized picture of justification that excludes cognitively unsophisticated agents like children or animals⁸, and some of its most popular versions are often accused of facing an uncomfortable regress problem⁹. Likewise, some well-known challenges for externalism include the Generality Problem¹⁰, the notorious issues surrounding cases involving accidentally reliable belief-forming methods (e.g. clairvoyance)¹¹, and the New Evil Demon Problem¹². The point is the following: whichever side of the internalism/externalism divide one stands on, the implications of one's decision are well explored. To put it differently, in the case of the internalism/externalism divide we have a pretty good understanding of (a) what's at stake and (b) what kind of objections one can expect to encounter depending on the position one sides with. In the case of the probabilistic/non-probabilistic divide however things are different.

Since people have only recently become interested in non-probabilistic accounts of justification, there has to this day been no systematic work exploring the implications that follow from one's decision to opt for either a probabilistic or a non-probabilistic account of justification. In other words, the implications of the divide are less explored and as a result it is a lot less clear what's at stake.

The first part of this study aims to address this problem by drawing out the major advantages and challenges for both probabilistic and non-probabilistic approaches to justification. In the end, I argue, that the probabilistic/non-probabilistic

⁷ See Goldman (1999) for a version of the challenge. For a response, see Conee & Feldman (2001).

⁸ See Alston (1988), BonJour and Sosa (2003).

⁹ For important discussions of this problem, see BonJour (1985, ch. 4), Alston (1988), and Bergmann (2006)

¹⁰ For an important discussion, see Conee and Feldman (1998).

¹¹ The famous case of Norman the clairvoyant was first introduced by Bonjour (1985). The case challenges the claim that reliability alone is *sufficient* for justification.

¹² The New Evil Demon Problem was first discussed in Lehrer & Cohen (1983) and Cohen (1984). Unlike Norman the clairvoyant, which challenges that reliability is *sufficient* for justification, the New Evil Demon Problem challenges the claim that reliability is *necessary* for justification.

divide leaves us in a dilemma, as neither approach is able to satisfy all the desiderata we would expect a wholly satisfactory theory of justification to accommodate.

The main objection against probabilistic theories of justification following the Lockean View, I argue, is that they are incompatible with multi premise closure – a compelling principle according to which we are justified in believing the conjunction of our individually justified beliefs.¹³

Multi Premise Closure (MPC) If S is justified in believing p and S is justified in believing q... and S is justified in believing n, then S is justified in believing the conjunction (p & q... n).

As we will see, there are a number of arguments supporting the claim that probabilistic accounts of justification are incompatible with (MPC). This is bad news for the Lockean View, as there are a number of interrelated reasons for thinking that a satisfactory theory of justification should accommodate (MPC). Here are two prominent ones.

First, (MPC) seems intuitively compelling - if someone is justified in believing a number of individual propositions, then surely they should also be justified in believing these propositions together. Furthermore, if we reject (MPC), then there would be nothing epistemically problematic about someone who justifiably believes each premise of an argument but refuses to believe the conclusion yielded by conjoining these beliefs. However, as we will see later, such cases do not just appear puzzling but also seems to indicate some rational failure on the side of the reasoner. For these, and other reasons, preserving (MPC) is often considered to be a plausible desideratum for a satisfactory theory of justification and the fact that probabilistic theories are unable to preserve it should count as a strike against them.

One of the primary motivations underwriting recent non-probabilistic accounts of justification is that unlike their probabilistic competitors they are able to

¹³ There are of course other objections and we will consider them in chapter 2, but, as we will see, proponents of the Lockean View have plenty of resources at their disposal to respond to these challenges.

preserve (MPC). However, as we will see, the Normalcy View faces a different problem. By requiring that in order to be justified a belief must be true in *all* normal worlds, normalcy views become too demanding. For instance, beliefs based on purely statistical evidence – like the belief that one is not going to win the lottery - regardless of how strong the statistical evidence, will always turn out to be unjustified on accounts following the Normalcy Views. This is the case because, as we will see, in cases where a belief that P is based on statistical evidence there always exists at least one normal world in which one falsely believes that P. Since many of us ordinarily do think that we can justifiably believe that we are not going to win the lottery, we may worry that normalcy views are somehow on the wrong track.¹⁴ More importantly however, I will argue the problem does not stop here - it is not just beliefs based on statistical evidence that turn out to be problematic for the Normalcy View. As we will see, normalcy views of justification give rise to accounts of epistemic defeat on which seemingly insignificant pieces of evidence turn out to have considerable defeating powers. This leads normalcy views to systematically deny the justificatory status of many beliefs we ordinarily take to be justified. This problem I argue exposes normalcy views to counterexamples and gives rise to serious skeptical worries.

Thus, the probabilistic/non-probabilistic divide, which may present one of the most fundamental divides in the justification debate, leaves us in a dilemma: If we opt for a probabilistic conception of justification, and hold that the property that makes a belief justified is the property of *being highly probable*, then we will not be able to preserve (MPC). Alternatively, if we opt for a stronger notion of justification, which can avoid this problem, and hold that the property that justifies belief is the property of being true in all normal worlds, then justification becomes too demanding and a lot of beliefs we ordinarily take to be justified will turn out to be unjustified. I take this dilemma – I call it the *(MPC)-Stinginess Dilemma* – to be one of the major obstacles we face in providing a satisfactory account of justification. So, how should we respond to this worry? Answering this question is the chief concerns of the second part of this study.

¹⁴ Aside from philosophers who have defended this position, there have also been a number of recent empirical studies supporting this claim; see Turri and Friedman (2014), Friedman and Turri (2015), Ebert et al. (2018).

3. *A Way Out of the Dilemma: The Functionalist Theory of Justification*

One immediate way out of the dilemma would be to simply accept one of its horns along with the unpalatable consequences that come with it. This strategy would have us simply bite the bullet. There are two options available.

Bite-the-Bullet-1 Accept the Lockean View and deny (MPC).

Bite-the-Bullet-2 Accept the Normalcy View and make do with a stingy account of justification.

However, I think it is easy to see that that neither of these options is wholly satisfactory. After all, both (MPC) as well as having a theory that is sufficiently generous seem like plausible constraints on a satisfactory theory of justification. And even if one thinks that the costs of either *Bite-the-Bullet-1* or *Bite-the-Bullet-2* are ultimately manageable, it is easy to see that a solution which avoids them would be *prima facie* more attractive. In short, it is conceivable that we could do better.

I am going to explore and defend an alternative way out of the *(MPC)-Stinginess Dilemma*, one which promises to avoid the costs associated with both *Bite-the-Bullet-1* and *Bite-the-Bullet-2*. The solution I propose promises to retain a version of (MPC) whilst avoiding the problem of being overly stingy. What makes this attractive mix of features possible? The key to the new solution is the rejection of a largely unchallenged assumption about the nature of justification, namely that there is only one way for a belief to be justified; or, to put matters differently, that there is only one property in virtue of which a belief can be justified. I will call this assumed monism about justification, *J-monism*.

J-monism There is only one way for a belief to be justified – i.e. there is only one property that can make a belief justified

The dilemma relies on this assumption since it is only on a monist picture of justification that we are forced to choose between a probabilistic account of justification - on which (MPC) is lost - and a non-probabilistic account, which will be too stingy. But this assumption is of course optional and can be rejected. A rejection of J-monism yields pluralism about justification - the idea that there is more than one way for a belief to be justified, or, more precisely, that there is more than one property that can make a belief justified. I am going to call this pluralist position *J-pluralism*.

J-pluralism There is more than one way for a belief to be justified – i.e. there is more than one property that can make a belief justified.¹⁵

Unsurprisingly, a pluralist theory of justification yields new resources for responding to old problems. One of the most important upshots of the pluralist picture is that the Lockean views and normalcy views of justification are no longer mutually exclusive but can instead work side-by-side to complement each other. For instance, J-pluralism is compatible with the idea that some beliefs are justified in virtue of *being highly probable* on the evidence and other beliefs are justified in virtue of *being normally true*. On this picture (MPC) would be preserved for all beliefs that are justified on the Normalcy View, while those beliefs that fail to be justified on the Normalcy View may still be justified according to the Lockean View. As a result, J-pluralism offers the promise of retaining a version of (MPC) whilst avoiding the problem of being too demanding. So, it appears that on a pluralist picture we can have our cake and eat it too.

The trick to pull off is to provide a compelling story for how J-pluralism may be turned into a workable theory of justification. After all, on its own, J-pluralism is simply a thesis about the nature of justification – by itself it does not do much to help

¹⁵ Let me highlight from the outset that the proposal ‘that there is more than one way for a belief to be justified’ is *not* supposed to be understood in the following trivial sense: some beliefs are justified by perceptual experience, some beliefs are justified by testimony, and yet other belief are justified by memory, so there is more than one way for a belief it be justified. If this is the pluralism I had in mind, then, since this is a rather uncontroversial proposal, it would turn out that most of us already are pluralists about justification. However, as mentioned above, the pluralism I have in mind is deeper and less trivial; when I say that there is ‘more than one way for a belief to be justified’ what I have in mind is that there exists more than one justification-conferring property – i.e. there exists more than one property in virtue of which a belief can come to be justified.

us answers the original epistemic question ‘what justifies belief?’. What is needed to turn J-pluralism into a workable theory of justification is an overarching metaphysics. Such a metaphysical framework is needed to provide answers to important metaphysical questions like the following: How do the different ways of being justified hang together – i.e. what is the relationship between them? How do we determine which notion of justification applies when? Does there exist some unifying property that all justified beliefs have in common or not – and if so what is this property? How do the different ways of being justified interact – e.g. how do we reason with different notions of justification and what happens if we conjoin beliefs that are justified in different ways? Without an answer to these questions, any pluralist theory of justification is bound to be incomplete and subsequently unsatisfactory.

Fortunately, in addressing these issues we can draw on the lessons learned from previous pluralist projects in other areas of philosophy - most notably truth pluralism championed by Crispin Wright (1992, 1999) and Michael Lynch (2004, 2009). Broadly speaking, truth pluralists defend a truth analogue of J-pluralism, according to which there is more than one way for a belief (or a proposition) to be true. More specifically, they defend the view that in different domains of inquiry truth is realized by different properties – in some *domains of inquiry* (e.g. the external world) truth is realized by the property of corresponding to mind-independent facts, while in other *domains of inquiry* (e.g. morality and mathematics) truth is realized by the property of being coherent with other propositions. Importantly, the rise of alethic pluralism grew out of a general pessimism regarding the prospects of alethic monism. As Lynch puts it, “In recent decades, many philosophers have come to think that the monist’s quest for the nature of truth is a fool’s errand.” (Lynch 2009: 4) And a key thought underlying this project is the worry that perhaps the monist’s quest for the nature of *justification* too is a fool’s errand. By drawing inspiration from the work on truth pluralism, particularly that of Michael Lynch, I will develop and defend a novel pluralist theory of justification, which I call the *Functional Theory of Justification*.

This functional theory has two primary tenets. First, that the property of *being justified* is a functional property the functional role of which is given by certain platitudes about the nature of justification. And second, that the property of being justified can be realized (or manifested) in more than one way, or by more than one

property (cf. Lynch 2009: 3). More precisely, like the truth pluralists who hold that in different *domains of inquiry* truth is manifested by different properties, I will argue that in different *epistemic environments* justification is realized by different properties.

Finally, I argue that the *Functional Theory of Justification* does not just offer a way out of the *(MPC)-Stinginess Dilemma* faced by its monist competitors; it also provides interesting new resources for responding to a number of epistemic puzzles, including the lottery paradox¹⁶, the preface paradox¹⁷, as well as more recent puzzles surrounding various legal cases, such as Gatecrasher and Blue Bus cases.¹⁸ Overall then, the pluralist theory proposed here seems to offer an interesting and attractive new theory of epistemic justification.

4. *The Road Ahead*

In what follows I will briefly outline the road ahead. The plan is as follows.

Chapter 1 is primarily methodological in nature. One question that naturally arises at the beginning of philosophical inquiry is what we should require of a theory in order to count as a theory of the target notion rather than something else. In the case of justification this means that we need to answer the following question: how do we determine whether any proposed property is a property that makes a belief justified rather than something else? In addressing this question, I defend what I call *The Minimalist Conception of Justification* (inspired by Wright's (1992) *Minimalist Conception of Truth*), according to which all there is to being a justification-conferring property is to satisfy a set of core platitudes, which capture what is widely taken to be essential to the nature of justification. An important feature of this minimalist conception is that it offers a natural way into pluralism, as there is no reason to think that the core platitudes can only be satisfied by a single property.

¹⁶ The lottery paradox was first introduced by Kyburg (1961).

¹⁷ The preface paradox was first introduced by Makinson (1965).

¹⁸ For a good overview of these puzzles see, Buchak (2014), Smith (2016), and Littlejohn (forthcoming)

Chapter 2 introduces the Lockean View - the currently most prominent picture of justification in the literature - according to which a belief that P is justified for S iff P is highly probable for S. I then consider but ultimately reject a number of recent arguments that set out to undermine the Lockean View on the grounds that it is compatible with the idea that beliefs based on purely statistical evidence can be justified. I consider but ultimately reject the most prominent of these arguments.

Chapter 3 develops what I take to be the most pressing objection to the Lockean View, namely that it is incompatible with the attractive and intuitive compelling principle of Multi Premise Closure (MPC). In arguing against the compatibility of the Lockean View with (MPC), philosophers often rely on two epistemic paradoxes: the *Lottery Paradox* and the *Preface Paradox*. However, in the past a number of people have proposed various ways of solving the two paradoxes without denying (MPC). As a result, the lottery paradox and the preface paradox have been unable to make a decisive case against the claim that Lockeans cannot preserve (MPC). The primary objective of this chapter is to introduce a new closure-threatening paradox that makes a stronger case against the compatibility of the Lockean View with (MPC) than its two predecessors. I will dub this paradox the Paradox of the Pill. As we will see, there is little hope that Lockeans will be able to avoid the conclusion that their views are incompatible with (MPC). Seeing that the preservation of (MPC) is a plausible desideratum for a satisfactory theory of justification, we have a good reason to be unhappy with the Lockean picture of epistemic justification.

Chapter 4 explores the prospects of non-probabilistic alternatives to the Lockean Picture; particularly that of the Normalcy View of justification, according to which a belief that P is justified only if there does not exist a single normal world in which S falsely believes that P. What makes normalcy views initially attractive is that they offer the promise of preserving (MPC). However, I argue that normalcy views face a serious problem when it comes to the notion of defeat. Using the leading versions of the Normalcy View - Martin Smith's (2010; 2016) *Normic Support Account* and Jarrett Leplin's (2009) *Normic Reliabilism* - as initial examples I argue that on the Normalcy View justification is too easily defeated. I call this the *Easy-Defeat Problem*. As a result, these views systematically deny justification for many beliefs we ordinarily take to be justified. This does not just expose the views to

straightforward counterexamples but also leads to a skeptical threat. Finally, I show that the Easy-Defeat Problem is not just an isolated problem for the accounts of Smith and Leplin but that we can expect the problem to generalize to any Normalcy View of justification. Therefore, just as we have reason to be unhappy with the Lockean View, so too do we have reason for being unhappy with the Normalcy View.

Chapter 5 does two things. First, it draws our attention to an uncomfortable dilemma that the work in the previous chapters points towards. This dilemma, which I dub the *(MPC)-Stinginess Dilemma* can be stated as follows: if we opt for a probabilistic account of justification and follow the Lockean View, then we will not be able to preserve (MPC); if on the other hand we opt for a non-probabilistic account of justification like the Normalcy View, then justification becomes too demanding and a lot of beliefs we ordinarily take to be justified will turn out to be unjustified. The remainder of the chapter develops a pluralist theory of epistemic justification – we will call it *The Functional Theory of Justification* – which unlike its monist counterparts is able to sidestep this dilemma. As we will see, the functional theory will be able to preserve a version of (MPC) whilst avoiding the problem of being too demanding.

Finally, in chapter 6 I argue that aside from helping us avoid the *(MPC)-Stinginess Dilemma*, the Functional Theory of Justification also provides compelling and interesting new solutions to a number of epistemic puzzles and paradoxes like the Lottery Paradox, the Preface Paradox, the Paradox of the Pill introduced in chapter 3, as well as a number of more recent puzzles involving statistical evidence. Importantly, due to the additional resources offered by the functional theory, it can deal with the puzzles in a way that avoids many of the costly and unattractive consequences faced by its monist competitors.

Chapter 1

Epistemic Justification: A Minimalist Conception

1. Justification and the Target of Inquiry

The chief concern of inquiry into the notion of epistemic justification is to provide an answer to the question, ‘What justifies belief?’. To make matters slightly more precise, we may say that the aim of inquiry into the notion of justification is to answer the question, ‘What property makes a belief justified?’. As such it is the job of a theory of epistemic justification to provide an answer (either partial or complete) to this question. In short, theories of justification are in the business of telling us what properties justify belief.

At this point however - before we have even started theorizing - we face a challenge. How do we determine whether any proposed theory of epistemic justification is really a theory of justification and not a theory of something else? Put in terms of properties we may state the worry as follows: how do we determine whether any proposed property is a property that makes a belief justified – or is a *justification-conferring* property - rather than something else? This issue is of course not an isolated problem for epistemic justification but a more general methodological problem for philosophical inquiry. Here is a plausible response to the problem: At the beginning of inquiry we need to identify some criteria that any theory needs to satisfy in order to count as a theory of the target notion rather than something else. But how might we go about establishing these criteria? And what are the relevant criteria in the case of epistemic justification? Addressing these initial methodological issues is the primary aim of this chapter.

In what follows I defend what I will dub *The Minimalist Conception of Justification*. According to the minimalist conception, there exists a number of core

platitudes, which capture what is widely believed to be essential to the nature of justification, and any property that satisfies the core platitudes presents a potential way in which a belief can be justified. One important upshot of this minimalist conception is that it is compatible with both *monism* about justification – the idea that there is only one way in which a belief can be justified – as well as *pluralism* – the idea that there is more than one way in which a belief can come to be justified.

2. *The Platitudes-First Approach*

When we inquire into the nature of truth, we set out to find a property that makes a belief true – so we are looking for a T-property; when we inquire into the nature of beauty, we set out to find a property that makes an object beautiful – so we are looking for a B-property; and when we inquire into the nature of justification, we set out to find a property that makes a belief justified – so we are looking for a J-property. More generally then we can say that, when we inquire into the nature of some target notion F, we set out to find a property that constitutes F – we are looking for an *F-property*. However, as mentioned above, before we can begin theorizing about the nature of some target notion - e.g. truth, beauty, or indeed justification – we first need some idea as to what we are looking for. This is important for at least two reasons. First, setting out to find something without knowing what we are looking for is difficult. And second, we need some way of evaluating whether any candidate theory of F is really a theory about F and not some other thing G. So, to get philosophical theorizing off the ground and to make sure that our theorizing stays at least broadly on track, we need to specify a set of criteria to guide our inquiry. Since they are *inquiry guiding* criteria, they should of course not just capture *any* feature associated with the target notion, but they must somehow capture what is thought to be *essential* to it. Borrowing some Lockean terminology from Lynch (2009: 7), we might say that the criteria should capture the target notion's *nominal essence* (cf. Lynch 2009: 7).¹⁹

¹⁹ For our current purposes we can think about something's *nominal essence* as follows: A specification of F's nominal essence is simply a specification of our ordinary notion of F.

How might we go about determining for any philosophical (or non-philosophical) notion what its essential features are? One natural *starting point* for such a project is to consult our ordinary thinking about the target notion. The idea goes as follows: by reflecting on the target notion we can begin to assemble a list of principles that capture the essential features of the concept we are interested in.²⁰ Following the terminology used by, for instance, Wright (1992, 1999) and later Lynch (2009), we may call these basic principles that capture what is widely believed to be essential to the nature of some target notion its *core platitudes*. Before we start theorizing about some target notion F then, we first need to compile a list of core platitudes that capture what are widely believed to be essential features of F.

Following the approach outlined above, the list of core platitudes for some notion F should be populated by what is widely taken to be essential to F. Borrowing again from Wright we might say that our list of platitudes should initially contain those things that chime well with our ordinary thinking about the target notion (1999: 226). So, by telling us what is widely taken to be essential to the target notion F, the core platitudes provide criteria that tell us whether some property is a candidate F-property or not. For instance, in the case of justification, justification's core platitudes will tell us what is required of a property to be a potential justification-conferring property – i.e. to be a potential J-property. These considerations support the following necessity claim.

Platitudes-Nec For any property G and target notion F, G is a candidate F-property *only if* G satisfies F's core platitudes.

From Platitudes-Nec it follows that if some property fails to satisfy F's core platitudes, then it cannot be a candidate property for constituting F. Again, we can consider justification as an example; if some property fails to satisfy justification's core

²⁰ In his work on truth Wright puts this idea as follows, "Faced, then, with the manifest improbability of an illuminating necessary-and-sufficient-conditions analysis of truth, there is still a more relaxed program of analysis....This more relaxed project will see us trying to build an overall picture of the concept of truth – of its contents and purposes – by the assembly and integration of as wide a variety as possible of basic a priori principles about it – platitudes, as I've elsewhere termed them." (1999: 226).

platitudes, then it is not a candidate justification-conferring property. It may of course still be an interesting property, but it is not a property that justifies belief. Platitudes-Nec allows us to answer an important question raised at the beginning of the chapter, namely how do we determine whether a theory is a theory of F rather than a theory of something else? According to Platitudes-Nec, for a theory to be a theory of F it needs to satisfy F's core platitudes.

There is of course room for disagreement about the core platitudes - after all the core platitudes, on the story I offered, are just based on our ordinary beliefs about the target notion's essential nature. For any target notion F, it may of course turn out that we were widely confused, unclear, or mistaken about its nature and that there are good reasons for rejecting some principles' alleged status of a core platitude. Some may reject what others take to be a core platitude.

However, since the core platitudes capture what is widely taken to be essential to the target notion, rejecting any of the core platitudes will lead to a number of challenges. First, proponents of theories about F that do not capture what many take to be essential to F will need to offer an explanation as to why people should have been so widely mistaken about the nature of the target notion. Additionally, besides offering an explanation of why people have been widely mistaken about the nature of F, they also need to provide a compelling reason for thinking they are in a somehow more privileged position than the rest of us when it comes to questions surrounding the nature of F. After all, we might think that if all of us have been so widely mistaken about the nature of F, then surely it is quite likely that others too will get things wrong. Second, theories that fail to capture the target notion's core platitudes will, as a result of rejecting what many take to be essential to the notion, strike many as implausible or as having implausible implications. Proponents of these theories will of course need to do some extra work to make palatable a theory that strikes a large number of people as implausible. Finally, any theory of F that fails to capture what many take to be essential features of F, will immediately invite disagreement about whether the offered theory really is a theory about F – the very challenge that we were hoping to address by identifying a list of relatively uncontroversial core platitudes. So, while it is possible to reject one or more of the core platitudes, in doing so one will incur a considerable explanatory burden.

Michael Lynch succinctly summarizes this platitude centred approach to philosophical inquiry outlined above as follows.

“... [w]hen setting off to discover the nature of some target property it helps to have some understanding of what it is we are looking for: its *nominal essence*, as Locke might have put it. The nominal essence of F, in the sense I intend, is our folk concept of F. It embodies our preconceptions, the way we tacitly think about it in ordinary life – even if, normally speaking, we don’t even recognize ourselves as doing so. A natural way of identifying something’s nominal essence, therefore, is to appeal to the set of largely implicit beliefs we folk have about it. By appealing to those folk beliefs, or truisms, we won’t typically learn *everything* about the object property we are interested in. And our later discoveries may force us to revise our preconceptions of it... But however these questions play out, keeping an eye on our folk beliefs about the thing about which we are curious will hopefully tell us whether our subsequent theories of its nature address the topic we were concerned with when our theorizing began.” (2009: 7-8)

So far, we have argued that in order for a property to be a potential F-property, the property needs to satisfy F’s core platitudes. A question that may arise at this point is whether being a candidate F-property requires more than satisfying the central platitudes. In other words, we have argued for Platitudes-Nec but should we also accept Platitudes-Suff?

Platitudes-Suff For any property G and target notion F, G is a candidate F-property *if* G satisfies F’s core platitudes.

I think there are good reasons for being sympathetic to Platitudes-Suff. Recall that the core platitudes capture what is widely taken to be *essential* to the target notion. And what more could we require of a candidate F-property than it satisfying what we take to be essential to F? By capturing F’s core platitudes, the property satisfies all the features that are thought to be essential to F. Nothing more seems to be required to

recognize a property as a candidate F-property. Paraphrasing from Wright (1999: 228) we might say that there is nothing essential to our concept of any target notion (in his case truth) that is not captured by the core platitudes, and thus no more to a theory being a theory of the target notion than it furnishing a model of the core platitudes. Consider again the example of justification. We are looking for a property that makes a belief justified; what more could we require of a candidate property than it being able to satisfy the things we take to be essential to the notion of justification? After all the property satisfies all the features that are widely thought to be essential to the notion of epistemic justification.

There is of course no guarantee that any theory, only because it satisfies F's core platitudes, will turn out to be a *good* theory, but we should nevertheless accept that it is *a* theory of F – even if it is ultimately unsuccessful. We should not conflate something *being a theory of F* with something *being a good (or persuasive) theory of F*. It seems compelling to think then, that satisfying F's core platitudes is sufficient for a theory to be a theory of F. Taken together Platitudes-Nec and Platitudes-Suff, make up what we may call the *Platitudes-First-Approach* to philosophical inquiry.

Platitudes For any property G and target notion F, G is a candidate F-property *iff* G satisfies F's core platitudes.

According to the Platitudes-First-Approach to philosophical inquiry, a notion's core platitudes take centre stage. And the aim of philosophical inquiry then is to find properties that can satisfy the core platitudes of the notion we are interested in. In other words, philosophical questions like 'What is truth?', 'What is beauty?', or 'What is justification?' essentially reduce to questions of the following kind: 'What property satisfies truth's core platitudes?', 'What property satisfies beauty's core platitudes?', or 'What property satisfies justification's core platitudes?'. Importantly, according to Platitudes, any property that does satisfy F's core platitudes qualifies as a potential F-property.

One important consequence of the *Platitudes-First Approach* is that it is compatible with both *monism* as well as *pluralism* about the target notion. Since any property that satisfies F's core platitude is a potential F-property and since there may

well be more than one property that satisfies the core platitudes, there is nothing that prevents the possibility of there being multiple F-properties. As Wright puts it, “there is no reason to expect that the minimal platitudes will constrain their interpretation to within uniqueness.” (1992: 75).

According to Platitudes First-Approach then, all there is to being a property for making beliefs (or propositions) true is to satisfy truth’s core platitudes; all there is to being a property that constitutes beauty is to satisfy beauty’s core platitudes; and all there is to being a property that justifies belief is to satisfy justification’s core platitudes. In the end, there may only be one property that satisfies the respective sets of core platitudes or there may be many. In this sense, the Platitudes-First-Approach yields what we may call *minimalist conceptions* of the notions we are interested in. And since the minimalist conception is perfectly compatible with pluralism, there is no reason to think that there will necessarily be a *single privileged property* that makes a belief true, constitutes beauty, or that makes a belief justified – it may turn out that these properties can be realized in more than one way.

What does this mean for our inquiry in to the notion of epistemic justification? It means that before we start theorizing about justification, we first need to identify some core platitudes that capture what we somehow take to be essential to the nature of epistemic justification and then see what properties can satisfy these platitudes. It also means that we should not be surprised to find that in the end there is more than one property that satisfies the core platitudes. In the next section I follow the Platitudes First-Approach outlined above to yield the *Minimalist Conception of Justification*. Doing so will of course involve the undeniably difficult task of identifying a number of core platitudes that capture what many take to be essential features of epistemic justification.

3. *Platitudes about Justification and the Minimalist Conception*

Coming up with a list of widely endorsed features about the nature of any philosophical notion is never an easy task. This is partially due to the fact that in philosophy there exists a considerable amount of disagreement about almost everything. And the notion of justification is of course no exception to this. While it

is important to acknowledge that a healthy amount of disagreement plays an important dialectical role in philosophical inquiry, which make it ultimately desirable, there is no denying that it makes the task of compiling a list of core platitudes more difficult. Generally speaking, core platitudes – given that they are supposed to capture what is essential to the target notion - should not be subject to widespread disagreement. In fact, if a principle is subject to widespread disagreement, then this may be taken as evidence that it is not a core platitude – after all it is not widely thought to be essential to the target notion. Looking at the notion of epistemic justification, for instance, the idea that justification is an exclusively internal matter – an idea which not too long ago was almost universally accepted – is now rejected by many and therefore should not be considered a platitude. Likewise, the idea that justification is closed under multi premise deduction, which was recently proposed as an essential feature of justification by Leplin (2009: 9)²¹, while intuitively compelling, is highly controversial and therefore should not feature in the list of justification's core platitudes. In short, in compiling a list of core platitudes for any philosophical notion, we should refrain from including anything that is subject to widespread disagreement. So, what are some core platitudes about justification?

3.1. Fallibility, Independence, and Truth Candidacy

Many people accept the claim that beliefs in some way or other 'aim at truth'. The phrase was coined by Williams (1973: 148) and while there is some disagreement about how to best spell out this metaphor, it is generally agreed that epistemically speaking our beliefs are truth-directed; or, put differently, that we aim to believe

²¹ He calls it an 'adequacy condition'.

truly.^{22,23} The widely accepted truth-directedness of belief serves as a helpful starting point for thinking about what's essential to justification. But what exactly does this connection between justification and truth look like?

One, almost universally endorsed, feature of the relationship between justification and truth is that justification is *fallible* – i.e. that we can have justified false beliefs. In other words, it is almost universally accepted that in order for belief to be justified it is not *necessary* that the belief be true. This will be our first core platitude about justification.

Fallibility In order for a belief to be justified it is not *necessary* that it be true.²⁴

²² The two most prominent interpretations are, using Chan's (2013: 3) terminology, a *normativist* interpretation and its *non-normative* opponents. The normativists (e.g. Wedgewood 2002, 2013; Shah 2003; Engel 2005, 2013; Shah and Velleman 2005) argue that beliefs' truth-directedness is a *constitutive feature* of beliefs. On this picture, aiming at truth is just a constitute part of what beliefs are. More specifically, proponents of the normativist interpretation argue that beliefs are governed by a *constitutive truth-norm* along the following lines.

- (1a) It is correct to believe that p if it is true that p; and
- (1b) It is correct to believe that p only if it is true that p. (Chan 2013: 4)

Opponents of the normativist interpretation of the metaphor (e.g. Velleman 2000; Horwich 2013; Papineau 2013) reject the idea that beliefs are governed by a truth norm – partially because they are skeptical about the idea that beliefs should have any *constitutive norms*. Instead, they argue that spelling out the metaphor requires appealing to factors external to the concept of belief. They usually explain the truth-directedness of belief *teleologically* by arguing that believing truly is in some way *desirable* or *valuable*. Papineau, for instance, spells out the metaphor by appeal to the biological function of belief, which is supposed to provide an explanation for why beliefs are truth-directed. For a good overview of the debate, see Chan (2013).

²³ Some writers use the observation that beliefs are truth-directed to derive highly ambitious, more general, epistemic goals. For instance, Alston has famously defended the view that our epistemic goal is to, “amass a large body of beliefs with a favorable truth-falsity ratio” (1985: 59). Similarly, Latus (2000) has claimed that our ideal epistemic goal is, “...to believe all the truths there are and only those truths” (31). Seeing that this goal is practically unachievable Latus proposes the following as our *intermediate* epistemic goal: “Our epistemic goal is to amass a large body of interesting true beliefs without also holding too large a body of false beliefs.”(31) I want to remain neutral on whether our epistemic aim is really is as ambitious and grandiose as Alston and Latus suggest. For our current purposes we do not need to follow them in adopting such an ambitious epistemic goal.

²⁴ Fallibility yields the result that theories according to which justification is factive (e.g. Sutton 2007; Littlejohn 2012) fail to be theories of justification. Instead they are better thought of as theories of some other notion distinct from epistemic justification. Is this a problem? I don't think so. First, there are very few of these views – perhaps precisely because they seem to involve the denial of what is widely perceived to be essential to justification. And second the suspicion that these theories miss out on essential features of justification is widely echoed in the literature, see Coffman (2010), Kelp (2011), Baumann (2014), Smith (2014), Grundmann (2015). In the end, not any theory, only because it is called a theory of justification by their proponents, needs to count as a theory of epistemic justification.

The key motivation for Fallibility is the idea that sometimes we have really good reasons for believing some proposition – e.g. because the evidence strongly favours the proposition or because we formed the belief using a highly reliable method – but, due to unfortunate circumstances, our belief turns out to be false. In these cases, many would want to say that we were nevertheless justified in holding the belief; after all, it was highly likely to be true.

A second, closely related, principle about the relationship between justification and truth that chimes well with our ordinary thinking is that not all true beliefs are automatically justified - i.e. that truth and justification are *independent*.

Independence Truth is not *sufficient* for justification.

Here is the thought that motivates Independence: if someone truly believes that P despite the fact that the evidence in support of P is very weak or the belief was formed by a highly unreliable belief-forming method, then this is not enough to make their belief justified. In order for a belief to be justified, it requires more than mere truth.

According to Fallibility and Independence then, truth is neither necessary nor sufficient for justification. So, what is the relationship between justification and truth? Based on the above considerations it seems compelling to think that justification is somehow concerned with those beliefs that, epistemically speaking, are particularly promising candidates for achieving our epistemic goal of believing truly. In other words, justification seems to be a property instantiated by those beliefs that are, from an epistemic perspective, particularly *good candidates* for being true.²⁵ This then will be our central platitude about the relationship between epistemic justification and truth.

Truth Candidacy (TC) The belief that P is justified for S only if, for S, the belief that P is a good candidate for being a true belief.

²⁵ While this idea is often simply taken for granted, it is explicitly defended in Bonjour (1985: 8), Alston (1985: 59), Latus (2000: 38), Douven (2008: 208), and Graham (2011: 57).

Now, there are of course some who disagree that beliefs aim at truth. Williamson (2000), as part of his broader knowledge-first project, has famously defended the *Knowledge Norm of Belief*, according to which one should believe that P, only if one knows P (11). Following Williamson's work some people have recently defended views on which both belief and justification are somehow 'aimed at knowledge' rather than truth. For instance, Sutton (2007) has defended the radical view that justification simply is knowledge, i.e. a belief is justified iff it is known. Bird (2007), Smith (2010, 2016), and Reynolds (2013) have endorsed weaker versions of this view on which justification amounts to something along the lines of 'would be knowledge'.²⁶ Finally, Ichikawa (2014) has defended the view that justification is *potential knowledge*. Most of these views (all but Sutton's), generally speaking, would endorse the view that a belief that P is justified for S only if for S, the belief that P is a good candidate for being knowledge. Let's call this idea, *Knowledge Candidacy* (KC). Does the recent popularity of (KC) pose a problem for (TC)'s status as a core platitude? Fortunately not. Since knowledge entails truth, anyone who endorses the view that in order for a belief to be justified it needs to be a good candidate for being knowledge, will also endorse that in order to be justified a belief will need to be a good candidate for truth – after all knowledge entails truth. In other words, a belief cannot be a good candidate for being knowledge without also being a good candidate for being true. So, since (TC) only expresses that truth candidacy is a *necessary* condition for justification it will also be accepted by anyone who accepts (KC).²⁷ Hence, (TC) is even more widely accepted than perhaps initially thought. This makes its status as one of justification's core platitude even more compelling.

Competing accounts of justification will of course tell different stories as to what makes a belief a 'good candidate' for being a true belief. In other words,

²⁶ The term 'would be knowledge' was first introduced by Sutton (2007: 10). But has also been used by Bird (2007) and McGlynn (2012). Following McGlynn (2012: 362) it is important to point out that Smith's motivation for endorsing the justification as 'would be knowledge' picture, does not seem to fall out of a more general commitment to the Knowledge-First program, which the other authors explicitly cite as their primary motivation.

²⁷ Those who endorse (KC) would of course reject that a belief being a good candidate for being a true belief is *sufficient* for justification, but (TC) does not make this claim.

competing accounts of justification will interpret the details of (TC) differently. For instance, traditional internalists would argue that epistemically speaking a belief is a good candidate for being a true belief as long as it is highly probable on the evidence²⁸. Alternatively, traditional reliabilists would argue that a belief is a good candidate for being a true belief given that it was produced by a generally reliable belief forming method. So, while internalists and externalists disagree about the perspective from which we should assess whether a belief is a good candidate for being a true belief, at a sufficient level of abstraction, many of them would still agree that what makes a belief a good candidate for being true is the property of being *highly probable*.

More recently a new family of views has emerged which provides a non-probabilistic interpretation of what makes a belief a *good candidate* for being a true belief. According to these views – we may call them *normalcy views* - the property that makes a belief a good candidate for being a true belief is the property of being true in all normal worlds (See Leplin 2009; Smith 2010, 2016).²⁹ What these considerations illustrate, is that the considerable degree of disagreement amongst epistemologists about the nature of epistemic justification can be understood, essentially, as disagreement about the correct answer to the question: what property makes a belief a good candidate for being a true belief? In other words, the disagreement is fundamentally about how to best spell out the details of (TC) – one of justification’s core platitudes.

Now, even though Fallibility, Independence, and Truth Candidacy are all core platitudes about justification – after all each of them captures a feature of justification that many take to be constitutive of the notion of epistemic justification – I am going to, in the interest of simplicity, treat Truth Candidacy as the primary platitude about the relationship between justification and truth, and will refer to the other two only when needed.³⁰

²⁸ Where evidence is construed internally.

²⁹ In chapter 4 we will look at these views in more detail.

³⁰ This of course does not mean that Truth Candidacy is more important than Fallibility and Independence. All the core platitudes are equally important. However, Truth Candidacy will be more relevant for the remainder of the discussion and so it will be helpful to prioritize it.

3.3 Permissibility

Another widely accepted feature about justification concerns its normative force. Since justification is a property of only those beliefs that, epistemically speaking, are particularly good candidates for being true beliefs, it is natural to think that justification also speaks to the question of what propositions we can permissibly believe.³¹ It seems compelling that if someone is justified in believing some proposition P, then it should be at least permissible for them to believe that P.

Permission (P) If the belief that P is justified for S, then it is *at least* permissible for S to believe that P.

There are of course some who think that epistemic justification offers something stronger than just a permission. For instance, Feldman (1988) and White (2005) have defended the view that if one is justified in believing P, then one has an *obligation* to believe P. However, even if one is inclined to accept this stronger view, one will also accept the claim that if an agent is justified in believing P, then it is *at least* permissible for them to believe that P. Permission then appears to capture a plausible way of thinking about the connection between justification and what we can permissibly believe.

3.4 Blamelessness

Finally, many believe that there is an important connection between the notion of justification and the notion of blame.³² Previously we considered the widely shared

³¹ This view has traditionally been endorsed by many of those who hold a ‘deontological conception of justification’, for instance, Bonjour (1985: 8), Goldman (1986: 59), Pollock & Cruz (1999: 11). But it also finds endorsement in more recent work; see Nelson (2010), Kroedel (2012). For a good overview, see Moser (1989) and Chuard & Southwood (2009).

³² While this idea is often taken for granted, it is more explicitly defended in, for instance, Ginet (1975: 28), and Bonjour (1985). Moreover, against the background of an increasing interest in the topic of epistemic norms, the notion of epistemic justification has recently been connected with the idea of *blameless* norm violation. Assume for a moment that truth is the Norm of Belief – i.e. that it is correct

idea that justification is fallible, i.e. that it is possible for a belief to be epistemically justified for S and yet turn out to be false. Related to the idea of fallibility, many accept that if a justified belief turns out to be false, then S is somehow free from blame. This popular idea captures the attractive thought that when someone has justification for believing P – e.g. because P is highly probable on the evidence or because it was formed using a highly reliable method - then the subject’s belief is, epistemically speaking, in good standing – after all P is a good candidate for being true. As a result, in cases of justified false beliefs, assuming the agent’s belief is based on the right reasons³³, many are reluctant to attribute blame to the agent; or, to consider the agent blameworthy. The following then is another plausible candidate for being a core platitude about justification.

Blameless (B) If S is justified in believing P, then, as long as S’ belief that P is based on the right reasons, S is blameless if P turns out to be false.

3.4 The Minimalist Conception of Justification

Together *Truth Candidacy*, *Permission*, and *Blameless* provide a set of core platitudes about the nature of epistemic justification. While I do not claim that this list is exhaustive – further reflection may yield additional features that are generally considered to be essential to the nature of epistemic justification - it should nevertheless serve as a good starting point for theorizing.³⁴ Following the Platitudes-First Approach, the core platitudes capture those features that are widely taken to be

to believe that P iff P is true. When a subject violates this norm – i.e. when they falsely believe that P - it may be suggested that a subject is blameless for this norm violation (or is excused for the violation) iff they justifiably believe that P. Here is how Brown (2017: 282), “One might suggest that one is blameless in violating a norm, N [e.g. the Truth Norm of Belief], if and only if one justifiably believes, or has justification to believe, that one’s not violating it (e.g. Arpaly 2002: 103-4).”

³³ This is to rule out cases in which a proposition is justified for a subject, but the subject does not form the belief based on the features that make the belief justified. In such cases, we may well want to judge an agent to be blame worthy for a false, yet justified, false belief.

³⁴ The thought that justification comes in degrees, or is *gradable*, might be another plausible candidate. But for our purposes here, we do not need to concern ourselves with this issue.

essential to the nature of epistemic justification - *Truth Candidacy, Permission, and Blamelessness* - and any property that satisfies these core platitudes can rightfully claim to be a theory of justification. In other words, the question of ‘What justifies belief?’ essentially reduces to the question of ‘What property can satisfy justification’s core platitudes - *Truth Candidacy, Permission, and Blamelessness?*’ and any property that can do so provides a potential way in which a belief can be justified. Let’s call this the *Minimalist Conception of Justification*. Why minimalist? Again, it is a minimalist conception because at its centre there is simply a list of core principles, which tell us what is required of a theory in order to count as a theory of justification; or, put in terms of properties, it is a minimalist conception because at its centre there is a set of core platitudes, which tell us what is required of a property to be a justification conferring property, or a J-property. Now, one important consequence of the minimalist conception is that it stays neutral on how many properties there are that satisfy the platitudes – i.e. it stays neutral on how many J-properties there will be. There may be just a single property that satisfies the core platitudes, in which case there will only be one way for a belief to be justified; or there may be multiple properties that satisfy the platitudes, in which case there will be more than one way for a belief to be justified.

4. *Advantages of the Platitudes-First Approach and Minimalist Conceptions*

So far, I have outlined a general methodological approach towards philosophical inquiry - *The Platitudes-First Approach* - and applied it to the notion of epistemic justification, which yielded the *Minimalist Conception of Justification*. In this section I am going to discuss some of the advantages of the Platitudes-First Approach and the minimalist conceptions it yields.

4.1. *Detecting and Avoiding Verbal Disputes*

In the past a number of people have argued that many debates in philosophy are merely *verbal disputes*.³⁵ Following Chalmers's definition, a dispute is merely verbal if "[an] apparent first-order disagreement arises wholly in virtue of metalinguistic disagreements." (2011: 525).³⁶ Put differently, in a verbal dispute parties take themselves to be having a substantive disagreement about F but really just *mean* different things by F. Here is a simple example. Person A uses the expression 'cold' to mean that the outside temperature is such that one should not go outside without putting on a coat first. Person B uses the term 'cold' to mean that temperatures are such that even when wearing appropriate winter clothes one will still be uncomfortable outside. Now we can easily imagine a disagreement between A and B, where B asserts that it is cold outside while A denies that it is cold outside. Their dispute however is purely verbal – A and B may well agree on all the first-order facts (e.g. the outside temperature, the level of humidity, etc.) and yet they are disagreeing – that's because their disagreement is really just a disagreement about the *meaning* of the term 'cold'. Now, those who think that verbal disputes are a widespread problem in philosophy, tend to think that many philosophical disagreements are in important respects similar to the above disagreement about the expression 'cold'. In other words, they worry that many philosophical disagreements are simply disputes about the meaning of certain expressions. What would verbal disputes in the philosophical domain look like?

Consider the following example. Imagine two people A and B who are disagreeing about whether some subject S is justified in believing that $2+2=4$. A uses the term 'justified' to apply only to beliefs about birds, while B uses the term to pick out beliefs that are particularly good candidates for being true beliefs (and also compatible with the other core platitudes). A keeps insisting that S is not justified in believing that $2+2=4$, while B is adamant that S is surely justified believing that $2+2=4$. Importantly, despite their disagreement, A and B agree on all the first-order facts (e.g. the degree of probability that S should rationally assign to the proposition

³⁵ For good discussions on this topic see Hirsch (2005) and Jenkins (2014).

³⁶ To be clear, this is the definition of *broad verbal disputes*, which take center stage in Chalmers' discussion. There is also a category of verbal disputes he calls *narrow verbal disputes*, but this distinction is not relevant for our current discussion.

that $2+2=4$, whether it is rational for S to act as if $2+2=4$, whether S knows that $2+2=4$, whether $2+2=4$ is true, whether $2+2=4$ is about birds etc.). Seeing that A and B seem to agree about all the relevant first-order issues, it seems that their disagreement is purely a disagreement about the *meaning* of the term ‘justified’. Hence, the disagreement between A and B seems to be a prime candidate for being a *purely verbal* dispute.

Why are some philosophers worried about verbal disputes? One hallmark of verbal disputes, as the examples above illustrate, is that they strike many as distinctively *pointless*. After all, in a verbal dispute, two parties are agreeing on all the first-order issues and are simply disagreeing about the *meaning of words*. Sometimes people use the metaphor of ‘talking past each other’ to characterize what’s going on in a verbal dispute. The worry of course is that if philosophers are frequently just talking past each other, then their disputes are unlikely to be philosophically fruitful or even interesting. In short, verbal disputes, at least on the surface, seem like pointless disputes that contribute little (if anything) towards the goal of philosophical progress. Let’s for the moment assume that those who worry that verbal disputes are a widespread phenomenon in philosophical discourse are right and that verbal disputes really are a serious issue in philosophy. In this case, detecting and avoiding verbal disputes would be an important philosophical matter. In what follows I argue that the Plitudes-First Approach significantly reduces the likelihood of purely verbal disputes. Moreover, it provides a straightforward method for detecting whether or not a dispute is verbal or not.

In an influential paper, Chalmers (2011) suggested that the very way philosophers traditionally ask philosophical questions makes them particularly liable to end up in verbal disputes (530-531). Philosophical questions often take the following form: ‘What is X?’

What is free will? What is knowledge? What is justification? What is justice?
What is law? What is confirmation? What is causation? What is color? What
is a concept? What is meaning? What is action? What is life? What is logic?
What is self-deception? What is group selection? What is science? What is
art? What is consciousness? And indeed: What is a verbal dispute?”

Despite their traditional centrality, disputes over questions like this are particularly liable to involve verbal disputes. (Chalmers 2011: 530-531)

Why are questions of this form particularly liable to verbal disputes? A natural way to answer a ‘what is X?’ question is to provide an account of what we *mean* by X. And with different people responding to questions about what they take certain philosophical notions to *mean*, we are bound to end up in verbal disputes. Why will the Platitudes First-Approach help us avoid verbal disputes?

On the Platitudes First-Approach, when inquiring into some notion X, we don’t ask the verbal-dispute-inviting question, ‘What is X?’. Instead, we first come up with a list of core platitudes, which capture what is essential to our ordinary notion of X, and any property that satisfies the core platitudes is a candidate X-property. Following this method, metalinguistic concerns, or questions about the *meaning* of certain expressions do not enter the picture. Sure, there may be disagreements about the core platitudes - i.e. about what is essential to the target notion - but these are substantive disputes over what is required in order for some theory to be a theory of the target notion rather than something else, they are not just disputes over the meaning of words. Hence, by avoiding the questions that are particularly likely to result in verbal disputes and by shifting the attention to more substantive issues – e.g. what the core platitudes are and what properties satisfy them - the *Platitudes First-Approach* will make it considerably less likely that we end up in merely verbal disputes.

To strengthen this claim, we might also point out that the Platitudes First-Approach comes close to the methodological framework that Chalmers’ argues will help us avoid verbal disputes. How does Chalmers think we can avoid verbal disputes? Since Chalmers’ thinks that it is questions of the form ‘What is X’ that get us into verbal disputes, it is not surprising that on his proposed solution we should start asking different questions. Here is what Chalmers has in mind.

“On the picture I favor, instead of asking “What is X?,” one should focus on the roles one wants X to play and see what can play that role. The roles in question here may in principle be properties of all sorts: so one focuses on the

properties one wants X to have and figures out what has those properties.
(2011: 538)

So, according to Chalmers instead of asking ‘What is X?’ we should focus on certain features we want X to have – for him these features concern X’s functional role – and then see what property can play these roles. The reason Chalmers thinks this methodological approach is preferable is because it stays clear of metalinguistic concerns about the *meaning of certain words*. Instead it shifts the focus to the target notion’s functional role. This of course comes close to what we said about the *Platitudes-First Approach*. Where on Chalmers picture we start out by first identifying a target notion’s functional roles and then see what property can play these roles, on the *Platitudes-First Approach* we start out by first identifying a set of core platitudes and then look for properties that satisfy the platitudes. Now, even though the *Platitudes-First Approach* and Chalmers’ proposed functionalism are not the same³⁷, they are nevertheless closely related. As a result, it seems reasonable to expect

³⁷ Chalmers’ Functional-Role-First Approach and the *Platitudes-First Approach* differ in a number of ways. First, the *Platitudes-First Approach* is not committed to functionalism. While it is of course perfectly *compatible* with a functionalist interpretation, on which the platitudes are taken to provide functional roles, this step is optional and by no means required. As such, Chalmers’ Functional-Role-First Approach may be seen as a more specific version of the more general *Platitudes-First Approach*.

Second, and perhaps more importantly, unlike the *Platitudes-First Approach*, Chalmers’ Functional-Role-First Approach does not provide a method for identifying a notion’s functional role. Recall that the *Platitudes-First Approach*, aside from avoiding verbal disputes, also helps us in answering questions about whether a theory is a theory about F rather than some other thing G. It is not clear that Chalmers’ Functional-Role-First Approach provides the resources required for answering these types of questions. Since on his favoured method we simply start by specifying a functional role and then see what property can play the functional role it is not clear how we would determine whether for any set of functional roles f, f describes the functional role of F rather than some other property G. For instance, in the case of epistemic justification, following Chalmers proposal we would simply specify some functional roles we want justification to play and then see what property can play these roles. But how would we determine whether the functional roles we have identified are those associated with justification and not some other property? The worry here is that Chalmers’ preferred picture might be *too uncommitted* or *too liberal* to answer these types of questions. This may provide a good reason for favouring the *Platitudes-First Approach* over Chalmers’ proposal.

These considerations point towards a third difference. Like the *Platitudes-First Approach*, Chalmers’ proposal is compatible with pluralism – after all for any set of functional roles it may be the case that there is more than one property that satisfies them. Let’s call this type of pluralism, *realizer pluralism*. However, Chalmers’ proposal, precisely because it stays quiet on how we determine a concepts’ functional role, also leads to a different type of pluralism – *conceptual pluralism*. Here is Chalmers, “This [his functional approach] leads naturally to a sort of conceptual pluralism: there are multiple interesting concepts (corresponding to multiple interesting roles) in the vicinity of philosophical terms such as ‘semantic’, ‘justified’, ‘free’, and not much of substance depends on which one goes with the term... I am inclined to think that pluralism should be the default view for most philosophical expressions.” (2011: 539). It is important to point out that this type of *conceptual*

that if Chalmers' *Functional-Role-First Proposal* is successful at reducing the risk of verbal disputes, then so will the *Platitudes-First Approach*.

The Platitudes-First Approach helps us avoid verbal disputes in another way: it provides an easy way of detecting instances of verbal disputes. Consider again that on the Platitudes-First Approach we start out by identifying a set of core platitudes, which capture what is widely believed to be essential to some target notion F. These core platitudes provide criteria that a property needs to satisfy in order to be candidate F-property. Importantly, if it does not satisfy the core platitudes, then it is not a candidate F-property. Now consider the above example of a purely verbal dispute, in which A and B disagree over whether S is justified in believing that $2+2=4$. A uses the term 'justified' to pick out only those beliefs that are about birds, while B uses the term to pick out only those beliefs that are good candidates for being true beliefs (and in a way that is compatible with justification's other core platitudes). Now the fact that A's usage of term does not satisfy any of justification's core platitudes while B's does, seems to provide a good reason for thinking that A and B mean different things by the term 'justified'. In other words, the fact that unlike B, A's usage of the term does not satisfy what is widely taken to be essential to the notion justification, provides a good reason for thinking that A and B are 'talking past each other'. Thus, by being explicit about principles that a property must satisfy in order to be a candidate property of the target notion, the Platitudes-First Approach provides a quick and easy way for determining when two (or more) parties are talking past each other, or, put differently, when the disagreement between two (interlocutors) seems to be best explained by the fact that they *mean* different things by same expression.

pluralism is very different from *realizer pluralism*. For instance, in the case of justification Chalmers' conceptual pluralism, which falls naturally out of his Functional-Role-First Approach, suggests that there may be *many different concepts of epistemic justification*. The Platitudes-First Approach is not committed to (and perhaps even incompatible with) this claim. While the core platitudes tell us what is widely taken to be essential to the notion of justification, and while it may be the case that there is more than one property that satisfies the platitudes, there is nothing that suggests that there is more than one concept of epistemic justification. As we can see, even though there are important similarities between Chalmers' proposal and the Platitudes-First Approach, there are nevertheless a number of differences. While I think that the Platitudes-First Approach presents an overall more satisfactory method for philosophical inquiry, none of the arguments in this study rely on this claim.

Thus, if we are worried about verbal disputes in philosophy, then the *Platitudes-First Approach* should be good news. As we have seen, the Platitudes-First Approach, by moving the attention away from the meaning of terms and towards more substantive issues considerably reduce the risk of philosophical inquiry leading to *purely verbal disputes*. By extension, we can expect that as we inquire into the notion of epistemic justification, the *Minimalist Conception of Justification*, since it is a product of the Platitudes-First Approach, will help us stay clear of merely verbal disputes about the term ‘justification’.³⁸

4.2. Locating and Diagnosing Genuine Disagreement.

Another closely related advantage of the Platitudes-First Approach to philosophical inquiry is that it will be helpful for locating and diagnosing genuine philosophical disagreements. In the case of justification, this means that the *Minimalist Conception of Justification* will provide a way of diagnosing and clarifying disagreements about the notion of epistemic justification. On the Platitudes-First Approach there are at least three types of disagreements that we are likely to encounter.

First, there may be disagreement about the core platitudes. As mentioned earlier, one or more parties may disagree over whether some principle is a core principle or not. In such cases, A proposes a set of principles that are allegedly core platitudes for some notion F, while B denies of at least one of the proposed principles that it really is a core platitude.

Second, parties may disagree over non-core principles. What would this look like? Two parties may agree about all the core principles that a theory needs to satisfy in order to be a theory of F rather than something else. But, they may disagree about

³⁸ Note that so far in this section we have assumed that verbal disputes really are a serious philosophical issue. While this strikes me as plausible, it is not uncontroversial. For instance, anyone who denies that our beliefs about, or associations with, expressions play a meaning-determining (or reference-fixing) role may deny that there are any verbal disputes. For instance, on an externalist picture of semantic content, where the meaning of expressions is fixed by external factors, there will be no ‘talking past each other’. However, even if one endorses such a *semantic externalist* account, on which verbal disputes may not arise, the Platitudes-First Approach and the minimalist theories it yields still provide a promising method for philosophical inquiry. Why? The Platitudes-First Approach helps us focus on what’s important, or widely thought to be essential, to the target of inquiry and it encourages us to explore what sort of properties may (or may not) be able to capture these features. These are virtues of the Platitudes-First method that do not rely on the existence of verbal disputes.

whether a *compelling* theory should also satisfy some further conditions that are not part of the core platitudes – i.e. that are not strictly speaking necessary for a theory to be a theory of F. We may call these additional principles that one feels inclined to think a compelling theory should satisfy, *desiderata*. On this picture, the core platitudes are widely accepted largely *uncontroversial* features that a theory needs to satisfy in order to count as a theory of F. The desiderata on the other hand are additional features that one thinks a *compelling* theory of F should satisfy; as such the desiderata can be more controversial and do not necessarily need not be widely accepted. With this distinction in mind we can easily imagine a case in which two people *agree* on the set of core platitudes a theory needs to satisfy in order to be a theory of F but *disagree* about whether a *compelling* theory of F should also satisfy some further desideratum D - i.e. some further principle that is not part of the core platitudes.

Finally, two (or more) parties may disagree over whether some property satisfies a set of platitudes (core or non-core). In such a disagreement A might hold that some property F does satisfy a set of platitudes f, while B holds that F does not satisfy one or more of the platitudes in f.

Importantly, unlike verbal disputes, I take all these types of disagreements to be substantive and philosophically interesting disagreements. Moreover, on the Platitudes-First Approach it is very easy to diagnose and explain philosophical disagreements. Let's use the historical standoff between internalists and process reliabilists as an example.

Here is an admittedly highly simplified overview of the debate. According to internalists, broadly construed, a belief that P is justified iff it has property of being highly probable on the agent's total body of evidence³⁹. Reliabilist on the other hand, again broadly construed, argue that beliefs are justified iff they have the property of being formed by a historically highly reliable belief forming process. In an effort to settle their disagreement over the question 'what property justifies belief?' internalists devised various alleged counter examples to reliabilist theories of justification – famous example include *Norman the Clairvoyant* and the *New Evil Demon Problem*

³⁹ Where evidence is internally construed.

- and argued that these counter examples provide reasons to *reject* reliabilist theories of justification.

Now, on the Minimalist Conception of Justification the move of simply rejecting reliabilism is no longer straightforwardly possible. Recall that on the minimalist conception a theory is a theory of justification rather than something else as long as it satisfies justification's core platitudes. And, assuming that both internalists and reliabilist theories can satisfy the three central platitudes identified above, internalists cannot so easily reject reliabilist theories of justification. So, how should we understand their disagreement?

Here is my preferred diagnosis.⁴⁰ The disagreement between internalists and externalists should be understood as a disagreement of the second kind outlined above, namely a disagreement about the status of some non-core principles or *desiderata*. While internalists cannot, and should not, deny that reliabilism *is* in fact a theory of justification – after all, reliabilism does satisfy all of justification's core platitudes – they *can* deny that reliabilism is a *compelling* theory of epistemic justification. Following the above mould for disagreements on the Platitudes-First Approach, this means that there will be some non-core principles (or desiderata) that internalists endorse and that reliabilism cannot satisfy. If this is the case, then internalists can argue that while reliabilism may well be a theory of justification, it nevertheless fails to be satisfactory because it does not satisfy some allegedly plausible desiderata about justification. What might these desiderata be that internalists think reliabilism cannot satisfy? Here are two candidate principles that internalists may argue a plausible theory of justification should satisfy: (i) justification should capture our ordinary judgments about certain cases (e.g. *Norman the Clairvoyant* and the *New Evil Demon*) and (ii) justification should be action guiding. Next, internalists may argue that since (i) and (ii) are plausible desiderata for a theory of justification and since reliabilism does not satisfy (i) and/or (ii), reliabilism fails to be a compelling theory of justification.

At this point externalists can of course reject that (i) and (ii) are in fact reasonable desiderata for a theory of justification; or, if they accept (i) and (ii), and

⁴⁰ There may well be others, but this one strikes me as particularly plausible.

they also accept that their view cannot accommodate (i) and (ii), then they can attempt to modify their view to make it compatible with the two desiderata. But note that at this point the disagreement is a substantive one and the precise locus of the disagreement between internalists and externalists is in plain view.

So, unlike Chalmers, who includes the internalism/externalism debate on the list of disputes that are purely verbal, the Minimalist Conception of Justification suggests that their disagreement is in fact substantive. Internalists and externalists are not talking ‘past each other’ or are merely disagreeing over the meaning of words, but instead internalists and externalists are disagreeing about what desiderata we should or shouldn’t expect the notion of justification to satisfy.

Thus, the Platitudes-First Approach to philosophical inquiry, besides avoiding verbal disputes, also helps us identify and diagnose substantial philosophical disagreements.

5. *On the Relation Between Platitudes and Theories*

It is important to note that the core platitudes themselves do not provide full theories of epistemic justification. The core platitudes only tell us what features a theory will need to have in order to count as a theory of justification. The core platitudes themselves do not tell us anything about what properties may *make a belief justified*; or, put differently, they do not tell us anything about how a belief’s justificatory status is realized. Hence, the platitudes fall short of being accounts of justification. Instead, the relationship between the core platitudes and theories of justification should be understood as follows: the core platitudes provide minimal constraints that tell us what is required for a theory to be a theory of justification, and particular theories of justification propose properties that can satisfy the platitudes and that are candidate properties for making a belief justified. In other words, we need the core platitudes to tell us what justification is, its nature or essence, and we need theories of justification to tell us what properties make a belief justified. (cf. Lynch 2009: 19).

In the next few chapters I consider two broad families of views. According to one, the property that makes a belief justified is that of *being highly probable*; according to the other, the property that makes a belief justified is the property of

being normally true. However, in the end, we have reason to be unhappy with both types of views, as both fail to capture some of the plausible and popular desiderata for satisfactory theories of justification. By drawing on the Minimalist Conception of Justification I then propose a pluralist theory of justification, according to which a belief can be justified in more than one way, or by more than one property. This pluralist theory, I argue, provides a compelling alternative to its monist competitors.

Chapter 2

The Orthodox Picture: The Lockean View and the Challenge from Statistical Evidence

1. *The Lockean View*

Many epistemologists have endorsed the idea that the correct answer to the question of ‘What justifies belief?’ will amount to something along the lines of *high probability*. How does this prominent idea translate into theories of justification? The thought that justification is best explained in terms of high probability naturally leads to what we may call threshold views of justification. According to these views, there exists some *epistemic good-making feature*, and a belief that P is justified in so far as P’s degree of probability, given the relevant good-making feature, is above some threshold value t required for justification. As we will see, this probabilistic way of thinking about justification is so widespread that it may rightfully be regarded as one of contemporary epistemology’s leading paradigms.

Now, as is to be expected, different accounts of justification will of course disagree about what exactly the epistemically relevant good-making feature is. For instance, internalists usually identify the total body of evidence⁴¹ available to an agent as the epistemic good-making feature that determines a belief’s justificatory status.

Threshold-Internalism S is justified in believing P iff S’s total evidence warrants a degree of confidence y in P, where $y \geq t$

⁴¹ Where evidence is internally construed.

Versions of this view are endorsed by, amongst others, Bonjour (1985: 6), Moser (1989: 42), Feldman & Conee (2004: 100), Foley (1992), and Sturgeon (2008). While the details can of course be developed in different ways, the key feature of Threshold-Internalism – the one that makes it internalist - is that a belief's epistemic status is *entirely* determined by S's evidence and the degree of confidence it warrants.

Externalists, while defending *structurally* similar views, have of course identified different epistemic good-making features. Externalists, unlike internalists, do not explain justification purely in terms of the evidence available to an agent but instead allow features external to the agent to enter the picture. According to one of the most popular externalist accounts of justification – *process reliabilism* - the feature that determines a belief's justificatory status is the general reliability of the belief-forming method that produced it.⁴²

Threshold-Externalism S is justified in believing P iff the belief that P was formed by a reliable belief-forming method – i.e. if it was formed by a method with a degree of reliability y is such that $y \geq t$.

More specific versions of this view have been defended by Goldman (1979), Plantinga (1993: 18), Bergmann (2006: ch.6), and Comesaña (2009). Again, While the details can be developed in different ways, the key feature of Threshold-Externalism – the one that makes it externalist - is that a belief's epistemic status is at least partially determined by factors external to the agent. For instance, facts about whether a belief-forming method is reliable are not accessible to the agent – they are entirely beyond S's epistemic reach.

Now, as we have seen in the last chapter, it is almost universally agreed that justification is fallible – i.e. that it is possible for agents to have justified false beliefs. In fact, fallibility was identified as one of justification's central platitudes. As a result, it should not surprise us that internalists usually do not require that in order for a belief that P to be justified the evidence must *entail* P. Instead, they hold that in order to be

⁴² The *locus classicus* here is Goldman (1979). What does it mean for a belief-forming method to be generally reliable? According Goldman, a belief-forming method is generally reliable if it produces a *favorable ratio of true to false beliefs*.

justified, the evidence needs to make the belief that P *sufficiently* probable. Similarly, externalists, do not require that in order for a belief that P to be justified, it must be formed using a perfectly reliable belief-forming method. Instead they usually require the belief forming method to be highly – but not necessarily perfectly – reliable. In short, internalists, in spelling out Threshold-Internalism, and externalist, in spelling out Threshold-Externalism, usually go in for a threshold value $t < 1$.⁴³

Importantly, then, many of the leading accounts of justification – both internalist as well as externalist - explicate justification in terms of high probability. For our current purposes, we may consider threshold views of this ilk instances of the following, more general, conception of epistemic justification, which we may call the *Lockean View*.⁴⁴

Lockean View The belief that P is justified for S iff P is highly probable for S

Note, that according to the Lockean View, the property that makes a belief justified - i.e. the *J-property* - is the property of being highly probable. Different instances of the Lockean View will of course tell different stories about what precisely this property is – e.g. the property of *being highly probable* on the evidence or the property of *being highly probable* given the relevant belief-forming method – but, ultimately it will be some property in the vicinity of *high probability*.

2. *The Lockean View and Justification's Platitudes*

Before considering one of the major challenges to the Lockean View, one might wonder whether the Lockean View satisfies justification's core platitudes outlined in

⁴³ There is of course an obvious worry here: how do we determine the degree of probability required for justification in a non-arbitrary way? This is sometimes called the *threshold problem*. While it is an interesting objection to the Lockean View, we will not consider it here.

⁴⁴ The Lockean View is essentially a justification analogue of the Lockean Thesis, which explains rationality in terms of degrees of confidence.

Lockean Thesis It is rational for you to believe p just in case it is rational for you to have degree of confidence y in p, where $y > x$. (Foley 1992: 112)

the previous chapter – *Truth Candidacy, Permission, and Blamelessness*. How do we determine this? To determine whether a theory satisfies the platitudes, we need to check whether it can provide a plausible *explanation* for why the platitude holds (cf. Lynch 2009: 18). In other words, to satisfy the platitudes a theory of justification needs to be able to rationalize the platitudes. Let's look at each of the platitudes in turn.

Truth Candidacy. According to Truth Candidacy, a theory of justification must somehow accommodate the idea that a belief is justified *only if* it is a good candidate for being a true belief. Can Lockean provide a reasonable explanation for why their view makes sure that this is the case? It appears they can. Consider that on the Lockean View a belief needs to be highly probable in order to be justified. And since *being highly probable* is presumably something that qualifies a belief as a 'good candidate' for being a true belief, the Lockean View seems to have no problem accommodating Truth Candidacy.

Permission. Can proponents of the Lockean View explain why it is the case that if a belief is justified, then it is at least permissible to believe it? Again, it seems like they can. On the Lockean View any justified belief will be highly probable. And proponents of the Lockean View will of course insist that it should be permissible for subjects to believe what is highly probable for them. So, it seems that the Lockean View can reasonably satisfy Permission.

Blamelessness. According to the final core platitude of justification, it needs to be the case that if a justified belief turns out to be false, then, as long as the subject formed the belief on the appropriate grounds, they will be blameless. Can the Lockean View offer a plausible explanation as to why this is the case? Consider that on the Lockean View, if a subject is justified in believing P, P will be highly probable for S. This means that if S's belief that P turns out to be false, something highly improbable must have occurred. This provides a reasonable explanation for why instances of justified false beliefs strike us as blameless - it seems harsh to blame an individual for the fact that something highly improbable has occurred. Since on a fallibilist picture is motivated precisely by the idea that there is always some risk of falsely believing P, keeping this risk small seems to be the best we can do as epistemic agents.

So, as we may have expected, the Lockean View appears to satisfy justification's core platitudes. As a result, it presents a candidate theory of justification and subsequently provides a possible response to the question, 'What justifies belief?'

However, recently the Lockean View has become the target of a series of objections. These objections centre around the Lockean View's interaction with statistical evidence.

3. *The Lockean View and Statistical Evidence*

A quick look at the types of evidence we rely on day-to-day when forming beliefs is enough to reveal that evidence comes in many forms and shapes: some beliefs are based on perception, some are based on memory, others are based on statistical evidence, and yet others are based on testimony. While for most of these types of evidence it is rather uncontroversial that they can justify beliefs, there is some controversy over whether purely statistical evidence - sometimes also called *naked*⁴⁵ or *bare*⁴⁶ *statistical evidence* - can justify belief.

One feature of the Lockean View is that it seems perfectly compatible with the idea that statistical evidence can justify beliefs

Statistical Evidence Statistical evidence - as long as it is sufficiently strong - is sufficient for justification.

Since proponents of the Lockean View tend to accept that justification (and rational belief) aims at truth⁴⁷, Statistical Evidence may initially strike us as unproblematic. After all, statistical evidence can be an excellent guide towards truth; it can be just as truth-conducive as other types of evidence (e.g. perception, memory, testimony, etc).

⁴⁵ The term 'naked statistical evidence' is often used in legal theory and in the discussion of legal cases. See, for instance, Redmayne (2008), Enoch et al. (2015).

⁴⁶ See Blome-Tillman (2015).

⁴⁷ For an explicit defense of this claim, see Velleman (2000), Wedgewood (2002), Shah (2003), Gibbard (2005), Engel (2013), Horwich (2013), and Papineau (2013)

And in some cases, it can be even more truth-conducive than other types of evidence, which can undeniably justify belief. Thus, for Lockeans, statistical evidence will simply enter into the probability calculus like any other piece of evidence. As such, Lockeans may be seen as proponents of a type of *anti-exceptionalism*⁴⁸ about statistical evidence, according to which statistical evidence is, in no interesting respect, different from other types of evidence.

However, recent years have seen the rise of a growing number of views that reject Statistical Evidence. According to these views, statistical evidence – regardless of its strength – can never justify belief. Examples include the various *Knowledge Views* of justification, which extend Williamson’s knowledge-first program to the notion of epistemic justification. Sutton (2005), for instance has defended the view that one is justified in believing P only if one knows that P. Others have defended weaker versions of the view according to which justification is a kind of ‘would-be’ knowledge⁴⁹ (Bird 2007; Smith 2016, Reynolds 2013) or ‘potential knowledge’ (Ichikawa 2014). Since statistical evidence alone does not put someone in a position to know, Knowledge Views deny that statistical evidence can justify beliefs. We can expect this result to generalize to any knowledge-first inspired view that takes as its starting point the idea that justification (or rational belief) aims at knowledge rather truth.

Another family of views that deny Statistical Evidence are *Normalcy Views* like the sort recently proposed by Leplin (2009) and Smith (2010, 2016). According to these views (we will look at them more closely in chapter 4), a belief that P is justified for S, only if there exists no normal world in which S falsely believes that P. Since for beliefs based on purely statistical evidence there always does exist a normal world in which one falsely believes that P, Smith and Leplin argue that it is a natural consequence of their views that statistical evidence cannot justify belief.

So, while Lockeans will treat statistical evidence on par with other types of evidence, proponents of Normalcy and Knowledge Views are committed to the idea that there is something epistemically deficient about statistical evidence.

⁴⁸ I thank Claire Field for suggesting this term.

⁴⁹ The term ‘would-be knowledge’ was coined by Sutton (2007: 10)

In light of this contention, it is not surprising that the topic of statistical evidence has become one of growing interest in recent debates. Since pre-theoretically many of us do think that beliefs based on strong statistical evidence can be justified, we may consider this a strike against Knowledge and Normalcy Views of justification. However, opponents of the Lockean View have recently attempted to turn this argument around by raising a number of objections to Statistical Evidence. These objections, if successful, would not just go a long way towards undermining the Lockean View – undeniably the leading picture of epistemic justification on the market – but they would also provide us with a perhaps compelling reason to opt for one of its more recent competitors (e.g. Normalcy and Knowledge Views). In short, much hangs on whether these arguments succeed.

This chapter considers the leading objections to Statistical Evidence and finds all of them wanting. As we will see, Lockeans and others who feel inclined to accept Statistical Evidence have more room for manoeuvring than their opponents seem to think. We will also consider a related, yet more direct argument, against the Lockean View, which relies on our comparative judgments about beliefs based on statistical evidence and beliefs based on other types of evidence. However, as we will see, this argument too is ultimately unpersuasive. The conclusion that we should draw from this is that the epistemic status of beliefs based on statistical evidence does not provide a persuasive reason to abandon the Lockean View.

4. *The Argument from Intuitions*

Amongst epistemologists critical of Statistical Evidence (and the Lockean View more generally), it has recently become popular to claim that it is *counterintuitive* to think that beliefs based on statistical evidence could be justified (e.g. Buchak 2014; Staffel 2016; Littlejohn forthcoming). They defend the following claim.

Intuition Many have the intuition that beliefs based on purely statistical evidence are not justified.

Is Intuition correct? If it were, then we have a *prima facie* reason to reject Statistical Evidence and subsequently the Lockean View. But I don't think that Intuition is particularly convincing. In fact, I think that it is false.

Let's consider the most prominent example of beliefs based on statistical evidence in the philosophical literature: beliefs about lottery tickets. Imagine you have just purchased a single ticket in a very large, fair, lottery (let's just assume that 1,000,000 tickets were sold). Since the belief is overwhelmingly probable, according to Statistical Evidence, you would be justified in believing that you are not going to win the lottery. Now, those who endorse Intuition will claim that this result is counterintuitive. But is it?

There are good reasons for thinking that the intuitions against justification from statistical evidence are not as widely endorsed as its proponents claim. There is a considerable number of people who *do* find it intuitive that statistical evidence should justify beliefs. Considering the case of beliefs about lotteries, Nelkin (2000: 375-376) notes that while many have the intuition that we cannot *know* beliefs based on statistical evidence, we do not tend to have the same intuition when it comes to the question of whether we can *rationally believe* a proposition based on statistical evidence.⁵⁰ In fact, Nelkin notes that denying the justificatory status of lottery beliefs tends to strike many as straightforwardly false. Similarly, (Comesaña 2009: 9) calls the claim that a subject in lottery cases is not justified in believing that their ticket is going to lose “unfashionable”. Finally, Ebert et al. (2018: 111) recently noted that those who endorse the claim that lottery beliefs are unjustified, “rarely describe it as “intuitive” or widely accepted, and almost universally regard it as something standing in need of substantial argument.”

Moreover, since Intuition makes an empirical claim about our intuitions it is also subject to empirical evaluation. Fortunately, our ordinary judgments about lottery beliefs have been the subject of a number of recent studies.⁵¹ These studies have consistently found that the majority of people *do* in fact judge lottery beliefs to be justified. Hence, there are good reasons to be suspicious about Intuition.

⁵⁰ This position is also defended by Hawthorne (2004: 8).

⁵¹ See Turri and Friedman (2014), Friedman and Turri (2015), and Ebert et al. (2018).

These considerations allow us to say two things. First, appeal to intuitions alone does not provide a persuasive case against Statistical Evidence – the evidence is at best inconclusive. And second, in so far as one takes seriously the empirical work about our judgments of lottery cases, we have good reasons to think that any view of justification that denies Statistical Evidence (e.g. Normalcy Views or Knowledge Views) will fail to be descriptively adequate – i.e. it will fail to reflect how we ordinarily think about epistemic justification.⁵² For our current purposes, however, I do not want to make too much of the second, more antagonistic, claim.

5. *The Argument From Practical Reasoning and Assertion*

A second, more theoretical, argument against Statistical Evidence proceeds from reflections on *practical reasoning* and *assertion*. Consider again the case of someone holding a ticket in a large lottery. Staffel (2016) has recently challenged the idea that one can justifiably believe that one's ticket is a loser on the grounds that it seems somehow inappropriate *to act as if* one has lost the lottery. For instance, it seems irrational to simply throw one's lottery ticket away. Here's Staffel, "if I can rationally believe that my ticket will lose, why buy it in the first place? Why not throw it away?" (2016: 1725-1726).

We can run an analogous argument for assertions. In cases where a belief that P is based on purely statistical evidence, it seems inappropriate to outright assert P. For instance, it is often noted that the purely statistical grounds in lotteries make it somehow improper for one to assert (without further qualifications) that one's ticket is a loser. Mirroring Staffel's question we may ask, if I can rationally believe that my ticket is going to lose, then why can't I assert it?

What's underwriting these two challenges is the more general idea that if one is justified in believing P, then it should also be epistemically permissible to rely on P in one's practical reasoning or to assert that P. More specifically, the principles underwriting these arguments can be stated as follows.

⁵² For a similar claim, see Turri & Friedman (2014: 32).

Practical Reasoning If S is justified in believing P, then it is epistemically permissible for S to act as if P.⁵³

Assertion If S is justified in believing P, then it is epistemically permissible for S to assert that P.

The details of Practical Reasoning and Assertion can of course be developed in different ways, but for our current purposes these general formulations will suffice. With the two principles in place, we can now reconstruct the argument against Statistical Evidence that Staffel and others seem to have in mind as follows:

Argument from Practical Reasoning and Assertion

- A1. A belief that P is epistemically justified for S only if it is epistemically permissible for S to act as if P and/or to assert that P
- A2. For beliefs based on statistical evidence it is epistemically impermissible for S to act as if P and/or to assert that P
- Ac. Therefore, beliefs based on statistical evidence are not epistemically justified for S.

How might proponents of Statistical Evidence respond to this challenge? There are at least two plausible options.

Option 1. Proponents of Statistical Evidence can accept Practical Reasoning and Assertion and explain the impropriety of asserting or acting on beliefs based on statistical evidence as resulting from conflicts with other *non-epistemic* norms— e.g. *prudential norms* and *norms of conversational pragmatics*.

This strategy accepts (A1) and accepts that it seems somehow improper to assert or act on beliefs based on statistical evidence but denies that this impropriety is of an epistemic kind. So, this strategy denies (A2). How might this strategy be motivated?

Consider again Staffel's challenge that when it comes to the belief that one's lottery ticket is a loser - regardless of the size of the lottery - it seems somehow

⁵³ A version of this principle has been defended by Fantl & McGrath (2002, 2009)

inappropriate to act as if one's ticket was a loser, i.e. to throw one's lottery ticket away. Now, it seems that this impropriety can plausibly be explained by appeal to *prudential norms* rather than epistemic ones. What makes it irrational to throw one's lottery ticket away is that one already has the ticket and with it a non-zero chance of winning the lottery. Forfeiting this opportunity by throwing the ticket away seems *imprudent* – one forgoes a potentially substantial pay-off for nothing.⁵⁴ This position is of course perfectly compatible with the idea that purchasing a lottery ticket in the first place is epistemically speaking unjustified or irrational – after all one is overwhelmingly likely to lose. But, once one holds a lottery ticket, it would surely be imprudent to throw it away.

Can we run an analogous argument for the case of assertions? And if so, what non-epistemic norms do we violate when we outright assert beliefs based on purely statistical evidence?

A number of people have previously suggested that the impropriety of asserting beliefs based on statistical evidence – e.g. lottery beliefs – can reasonably be explained as a violation of certain *pragmatic norms of conversation*. Weiner (2005), for instance, has offered a Gricean explanation. He argues that there is a norm, similar to Grice's Cooperative Principle (1989), according to which one's utterance must have some point. Asserting that one's lottery ticket is a loser, according to Weiner, is in some sense pointless, "It would be pointless for Sarah to tell Alice that her ticket hadn't won unless Sarah had some inside information not obviously available to Alice. Accordingly, when Sarah asserts that Alice's ticket did not win, she implicates that she has inside information." (2005: 232). So, according to Weiner

⁵⁴ At this point, proponents of Statistical Evidence may also point out that it is not always seem impermissible to act on beliefs based on purely statistical evidence. For instance, in cases where the practical stakes are very high, many of us would accept that it is entirely unproblematic to act on beliefs based on purely statistical evidence. Imagine the following lottery variant. You are at a party and have just picked up one of the 1000 sandwiches provided by the host. Before taking your first bite, a close friend of the host walks up to you and warns you that, for some undisclosed reason, the host has poisoned 1 of the 1000 sandwiches chosen at random. In this case many of us would judge that it seems perfectly appropriate not to eat the sandwich - i.e. to act as if one's sandwich was poisoned - despite the fact that the claim is highly improbable and is supported by purely statistical evidence. Hence, proponents of Statistical Evidence may conclude that reflecting on the permissibility of acting on beliefs based on purely statistical evidence doesn't provide any reason for thinking that statistical evidence is somehow epistemically deficient.

the impropriety of asserting that one's lottery ticket is a loser is best explained *not* by the fact that one lacks justification for believing it but by the fact that it violates a norm of conversational pragmatics.⁵⁵

Proponents of Statistical Evidence then have a plausible response to the *Argument from Practical Reasoning and Assertion*: They can accept Practical Reasoning and Assertion and explain our reluctance to act on or assert beliefs based on statistical evidence as a conflict between these *epistemic norms* and other *non-epistemic norms*. In short, the reason why you cannot throw your lottery ticket away or assert that you have lost the lottery is not because statistical evidence is insufficient to justify beliefs but because doing so would violate other non-epistemic –e.g. *prudential* or *pragmatic* - norms. However, this is only one possible response.

Option 2. Proponents of Statistical Evidence can also reject Practical Reasoning and Assertion by denying that actionability and/or assertability are necessary conditions for a belief to be justified.

In contrast to Option 1, this strategy accepts (A2) but rejects (A1). Proponents of this strategy can accept that it is epistemically impermissible to act on, or assert, beliefs based on purely statistical evidence, but deny that this tells us anything interesting about the nature of epistemic justification.

⁵⁵ Lackey (2007) recently presented a similar argument. According to Lackey, assertions of beliefs based on statistical evidence – like lottery beliefs - are inappropriate not because they are pointless, as Weiner suggests, but because they are *misleading*. Asserting that one's ticket is a loser, Lackey argues, is misleading because it suggests that one has inside information, "In particular, in cases where it is truly improper to assert lottery propositions, such assertions will frequently be impermissible precisely because they are misleading, where this misleading nature is both directly relevant to the purposes of the exchange in question and reasonably anticipated by the asserter." (Lackey 2007: 618) In more general terms, Lackey argues that the inappropriateness of asserting beliefs based on purely statistical evidence is best explained in light of their violating the following *Not Misleading Norm of Assertion*.

NMNA** S should assert that *p* in context C only if it is not reasonable for S to believe that the assertion that *p* will be misleading in C relative to the purpose of the exchange in question.

One worry about Lackey's strategy however is that the impropriety of asserting that one's lottery ticket is a loser persists, even when the misleading implicature is cancelled. For instance, consider the following assertion, 'Look, I don't have any inside information and don't know any more about the lottery than you do, but your ticket is a loser'. This assertion still sounds infelicitous despite the fact there is no longer anything misleading about it – it no longer violates NMNA**. This worry seems to give proponents of Statistical Evidence a reason to prefer Weiner's account.

Here's how this response may be motivated. Note that on the face of it, Staffel's observation that it seems inappropriate to act on beliefs based on statistical evidence, only tells us something about the *norms governing practical reasoning* – e.g. that it's inappropriate to act on beliefs if they are supported by purely statistical evidence. Likewise, the observation that it's inappropriate to assert beliefs based on statistical evidence initially, only tells us something about the *norms governing assertion* – e.g. that it is inappropriate to assert beliefs based on purely statistical evidence. But why should we accept that constraints on assertion and practical reasoning are also constraints on justification? In other words, why should we expect, as Practical Reasoning and Assertion suggest, that a belief can only be justified if it also satisfies the constraints on proper assertion and actionability? This move can reasonably be rejected for a number of reasons.

First, note that Practical Reasoning and Assertion commits us to what is sometimes called *impurism* about justification. According to impurism, whether S is justified in believing that P is not determined solely by truth-conducive factors –e.g. the degree of evidential support or a belief forming method's degree of probability – but it's also at least partially determined by pragmatic considerations – e.g. whether it's appropriate for S to rely on P in practical reasoning and whether it is appropriate to assert that P.⁵⁶ The idea that a belief's epistemic status should be affected by pragmatic – i.e. non-epistemic factors – can of course be resisted. As mentioned earlier, many of us (especially Lockeans) think that rational belief and justified belief aims at truth. On this picture – we may call it the *purist* picture - justification is concerned purely with believing truly or with believing accurately. Against the background of purism, we should be reluctant to tie justification to notions like assertion or practical reasoning, which are influenced by pragmatic factors that are not aimed at truth-conduciveness or do not advance our goal for being accurate epistemic agents. Proponents of this strategy may warn us that we should not muddle the epistemic water by introducing into it pragmatic concerns. In short, purism about justification offers proponents of Statistical Evidence a good reason to reject Practical

⁵⁶ For a defense of purism about justification, see Fantl & McGrath (2002). For a good critical discussion of impurism, see Brown (2013).

Reasoning and Assertion and subsequently (A1) of the *Argument from Practical Reasoning and Assertion*

A second reason to resist that the notions of assertion and practical reasoning should feature in our analyses of justification is that the former can reasonably be expected to be governed by more demanding norms than the latter. After all, unlike beliefs, actions and assertions are *social acts* –i.e. they are acts than can have a considerable effect on others. As such it should not surprise us that the norms governing assertion and practical reasoning are more stringent than the conditions required to justifiably believe a proposition. Put differently, there are reasons to think that justification and the two notions - practical reasoning and assertion - are not as closely connected as Staffel and others think. A similar position was recently defended by Whiting (2013).

Precisely because asserting is ‘external’, rather than ‘internal’, it is, if not necessarily a social act, then necessarily a potential social act. As a result, in evaluating an assertion, one might have to take into account the effects it might have on others, the expectations and needs of one’s interlocutors, the part that speech act might play in unfolding conversation, and so on. Evidently, all these considerations are foreign to the assessment of belief.... In light of these observations it should not come as a surprise that a situation in which one may believe that p might not be a situation in which one may assert that p. (187-188)

The thought that practical reasoning and assertion are (and should be) governed by more stringent norms than mere belief provides proponents of Statistical Evidence with another reason to reject Practical Reasoning and Assertion. Following this proposal, the apparent impropriety of acting on or asserting beliefs based on statistical evidence should simply be taken as data about the correct *Norms of Assertion* and *Norms of Practical Reasoning*.⁵⁷ There is no obvious reason for thinking that it is data about the correct accounts of epistemic justification.

⁵⁷ For instance, many endorse some version of the Knowledge Norm of Assertion (e.g. Williamson 2000, DeRose 2002, Hawthorne 2004)

Knowledge Norm of Assertion (KNA) It is appropriate to assert that P iff one knows that P.

So, proponents of Statistical Evidence have at least two ways of responding to the Argument from Practical Reasoning and Assertion – the first one rejects (A2), the second one rejects (A1). There are of course more options. For simplicity I have assumed that we treat Practical Reasoning and Assertion symmetrically, i.e. that if we accept one, then we will also accept the other and vice versa. This however is not mandatory. Perhaps some proponents of Statistical Evidence one would like to accept Practical Reasoning but reject Assertion (or the other way around). In these cases we can combine Option1 and Option2 to yield a more bespoke solution to the Argument from Practical Reasoning and Assertion.

6. *The Argument from Legal Cases*

The final argument against Statistical Evidence advances from reflections on certain legal cases. The most explicit and detailed version of this argument was recently proposed by Littlejohn (forthcoming).⁵⁸ Littlejohn invites us to consider the following case originally presented by Redmayne (2008).

Prisoners 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The

Since statistical evidence does not put us in a position to know, asserting beliefs based on purely statistical evidence would violate (KNA). Hence, (KNA), if one accepts it, provides a straightforward explanation of why it is inappropriate to assert beliefs based on statistical evidence.

Likewise, some have defended an analog knowledge norm for practical reasoning (e.g. Hawthorne & Stanley (2008), and Fantl & McGrath (2009))

Knowledge Norm of Practical Reasoning (KNPR) It is appropriate to rely on P in practical reasoning iff one knows that P

Again, (KNPR), if one accepts it, would explain why acting on beliefs based on statistical evidence strikes us as inappropriate. As such, (KNPR) would answer Staffel's challenge.

But Even if one does not endorse knowledge norms of assertion and/or practical reasoning and instead prefers a truth norm (see Weiner 2005) or a reasonable belief norm (see Douven 2006, Lackey 2007, Kvanvig 2009), one could always modify one's favorite norm in a way that would rule out beliefs based on statistical evidence. One straightforward way of doing so would be to build an anti-luck condition into the norm of assertion and practical reasoning.

⁵⁸ Less explicit precursors of this argument can be found in Smith (2010, 2016), Buchak (2013), and Staffel (2016).

100th prisoner played no role in the assault and could have done nothing to stop it. There is no further information that we can use to settle the question of any particular prisoner's involvement.

Suppose that after the event the prison officials decide to legally pursue a randomly chosen prisoner from the yard for their involvement in the attack. Should this prisoner be convicted? Since the evidence against the prisoner is purely statistical, many of us are inclined to say no. This judgment is also reflected in our legal practices - purely statistical evidence like the one in Prisoners would be insufficient to find the prisoner guilty.⁵⁹ It seems then, that many of us, including legal scholars, reject what Littlejohn (forthcoming) calls Punish.

Punish It is permissible to punish the defendant in Prisoner (and similar cases where the only evidence of guilt is statistical evidence).

So far so good. Next, consider the following epistemic analogue of Punish, which Littlejohn calls Believe.

Believe It is permissible to believe that the defendant is guilty in Prisoners (and similar cases where the only evidence of guilt is statistical evidence).

Now, what should we make of Believe? While the intuitive pull to reject Punish is undeniably very strong, it is less obvious whether there is anything objectionable about Believe.

⁵⁹ For important discussions concerning this and relevantly similar cases involving purely statistical evidence in court – e.g. the Blue Bus Case and the Gatecrasher Case - see Cohen (1977), Thomson (1986), Wassermann (1991) Colyvan et al. (2001), and Redmayne (2008). More recently these cases have been used as data by epistemologists and legal scholars who are trying to determine precisely what feature (or property) a body of evidence needs to have to suffice for legal convictions. For instance, Enoch et al. (2012, 2015) have argued that in order to suffice for legal convictions a body of evidence needs to satisfy a sensitivity condition (see Blome-Tillmann (2015) for a critique). In contrast, Smith (forthcoming) has recently argued that a body of evidence needs to normically support the proposition that the defendant is guilty. This debate however is not immediately relevant to the argument advanced in this paper, as it does not rest on any particular view about the legal standard of proof. Instead I am interested in the arguments that try to draw *epistemic conclusions* from our judgments about legal cases.

Proponents of Statistical Evidence – e.g. Lockeans – will of course accept Believe. Here is a quick and straightforward argument they may offer in support of it: Since many of us (particularly proponents of the Lockean View) accept that justification is truth-directed and since the belief that the defendant in Prisoners is guilty is an excellent candidate for being a true belief, one is justified in believing that the defendant is guilty. So, the position that proponents of Statistical Evidence will likely want to take towards Prisoners is to (a) accept Believe but (b) reject Punish.⁶⁰

At this stage however opponents of Statistical Evidence may raise the following objection: if it is permissible to believe that the defendant in Prisoners is guilty, then why shouldn't we punish the defendant? Here is how Littlejohn states the challenge, "If we can reasonably believe that the defendant in Prisoner is guilty, there is no principled objection to conviction" (forthcoming: Sect. 6) The core of this challenge against Statistical Evidence then is that it is not clear on what grounds proponents of Statistical Evidence may reject Punish if they think that it's permissible to believe that the defendant is guilty. In other words, we might worry that proponents of Statistical Evidence will struggle to square an acceptance of Believe with a rejection of Punish. Here is a more detailed version of what I will call *The Conviction Argument* against Statistical Evidence.

The Conviction Argument

- C1. If one accepts Statistical Evidence, then one is committed to accept Believe.
- C2. If one accepts Believe, then there is no principled objection to Punish in Prisoners.
- C3. If there is no principled objection to Punish in Prisoners, then one should accept Punish.
- Cs. Thus, if one accepts Statistical Evidence, then one should accept Punish
- C4. We should not accept Punish.
- Cc. Therefore, we should not accept Statistical Evidence

How may those sympathetic to Statistical Evidence respond to this objection?

⁶⁰ Alternatively, proponents of Statistical Evidence could of course try to argue that we should accept *both* Believe and Punish. However, as we will see, this is not a persuasive position.

One option would be to deny (C4) by arguing that against what many of us think, we should accept Punish. A possible argument for Punish, which Littlejohn considers as well, is that it may maximize expected value. If all 100 prisoners were convicted, then we know that 99 of the 100 convictions would be correct and entirely unproblematic. If on the other we convict none of the 100 prisoners, then we know that 99 guilty prisoners go unpunished. So, accepting Punish would commit us to 1 error, while rejecting Punish will lead to 99 errors. Depending on what values we assign to *falsely convicting the innocent* and *failing to convict the guilty*, we can, as long as we are free to change the number of prisoners, always set up a case in which Punish will maximize expected value. However, ultimately, I don't think that this strategy is a promising avenue for proponents of Statistical Evidence - our inclinations to reject Punish are too strong.

Alternatively, we could reject one of (C1) or (C3). However, both claims seem very compelling. If we take seriously the idea that statistical evidence can justify beliefs, then we should accept Believe. Likewise, it seems plausible that if there is no principled objection to Punish in Prisoners, then we should accept Punish. While others are invited to explore these avenues further, it strikes me as unlikely that a successful response to the *The Conviction Argument* will turn on a rejection of either (C1) or (C3).

If that's correct, then rejecting (C2) of the Conviction Argument seems to be the most promising strategy for defendants of Statistical Evidence. For this strategy to succeed, they will need to offer a *principled reason* to reject Punish that is compatible with their commitment to Believe. I turn to such a reason now.

6.1. Against Punish

On what grounds might we reject Punish? Opponents of Statistical Evidence can reject Punish on epistemic grounds. For instance, Littlejohn (forthcoming: Sect. 5.3) takes the following to be the strongest argument against Punish - let's call it *The Epistemic Argument against Punish*.

The Epistemic Argument against Punish

D1. It would be wrong to punish the defendant in Prisoners if we could not rationally believe the defendant to be guilty.

D2. Given the grounds in Prisoners, we could not rationally believe the defendant to be guilty.

Dc. Thus, it would be wrong to punish the defendant in Prisoners.

While proponents of Statistical Evidence can accept (D1) they will reject (D2). As a result, they cannot rely on The Epistemic Argument against Punish. Instead they will need to offer a different, *non-epistemic*, argument against Punish.⁶¹

In what follows I argue that Punish can reasonably be rejected on *moral grounds*. Convicting a defendant based on purely statistical evidence seems to be incompatible with our ideals of *justice and fairness*. If this argument is successful, then proponents of Statistical Evidence have a principled objection to Punish that is entirely independent from Believe – i.e. they will have a well-motivated reason to reject (C2) of *The Conviction Argument*.

To initially motivate the thought that Punish is morally problematic, it may be helpful to reflect on *why* we so vehemently object to Punish. Here is a natural thought: deciding the fate of a randomly chosen prisoner solely based on the fact that they were in the same prison yard as a group of individuals who committed a crime appears to be grossly unfair or unjust towards that individual.⁶² Even if it's overwhelmingly

⁶¹ At this stage some readers may feel inclined to raise the following objection to (C2): (C2) is unsustainable as a general principle, so we should reject it. To illustrate this, consider the following generalized version of (C2).

C2-general If it is permissible to believe that a defendant is guilty, then there is no principled objection to convict the defendant.

C2-general fails for obvious reasons. Sometimes evidence against a defendant is obtained through illegitimate means (e.g. illegal wire-tapping, coercion, etc.). In these cases, it may be permissible to believe that a defendant is guilty – after all the degree of evidential support for the defendant guilt is very high - but there are principled objections against convicting the defendant. In fact, courts would not find the defendant guilty if the evidence was obtained via illegitimate methods. This then seems to be a clear counterexample to C2-general. Is this enough to undermine (C2)? The answer is no. (C2) is specifically about Believe and Punish in Prisoners, in which there are no concerns about how the evidence was obtained. So, while there are good reasons to think C2-general is obviously false, these reasons do not apply to the more specific principle (C2).

⁶² It may be important to highlight that in Prisoners the one innocent prisoner presumably did not have the option of removing themselves from the situation or to intervene. In other cases of group crimes

probable that they were involved in the attack and we accept that it is rational to believe that the defendant was involved in the attack, it still seems unfair or unjust to determine the defendant's future in this way. Our reluctance to accept Punish then can reasonably be explained by our concern over the fair and just treatment of individuals.

The idea that legal convictions based on statistical evidence are morally problematic is also echoed in various places in the literature. For instance, Blome-Tillmann (2015: 103) recently offered the following assessment of the Gatecrasher Case (an analogue of Prisoners in which a person is randomly chosen from a crowd at a rodeo at which 70% of attendees are known to be gatecrashers and is sued for gatecrashing).

... courts will find for the defendant. And again, I take it that this accords rather well with our intuitions about *fairness* and *justice*.⁶³ It doesn't seem right to find John, who was randomly picked out in the arena, liable to pay damages just because 70% of attendees gatecrashed. After all, if such a case was allowed to succeed in court, the organizers of the rodeo could, in principle, win similar cases for every person in attendance at the rodeo, including the 300 people that paid the entrance fee.

Similar remarks can be found in Smith (2016: 37), who provides the following observation in response to the Blue-Bus Case (another analogue of Prisoners, in which the Blue-Bus Company is sued for their involvement in an accident based purely on the fact that the Blue-Bus Company operates 95% of the buses in the area).

Should I go around announcing that the bus involved was a Blue-Bus bus?
Should I take steps against the company – boycott their buses, picket their offices,

where these are live options, but an individual does not pursue them, both our intuitions and legal opinions may differ.

⁶³ My emphasis

etc? If my only evidence is that 95% of the buses in the area on the day in question were Blue-Bus buses, then to take such steps would surely be *unjust*.⁶⁴

It appears then that proponents of Statistical Evidence are able to run the following argument against Punish – I will call it *The Moral Argument against Punish*.

The Moral Argument against Punish

- E1. It would be wrong to punish the defendant in Prisoners if this would violate a plausible norm of justice/fairness.
- E2. Convicting a defendant in Prisoners, would violate a plausible norm of justice/fairness.
- Ec. Thus, it would be wrong to punish the defendant in Prisoners.

This argument provides a principled objection to Punish that is entirely compatible with, or independent from, Believe. As such, it offers proponents of Statistical Evidence the resources necessary to reject (C2) of the *Conviction Argument*.

What may this principle of justice or fairness be that punishing the defendant would violate? While it is likely that there is going to be more than one reasonable explanation as to what makes Punish unfair or unjust, and while the success of the argument does not require a commitment to any particular principle of justice or fairness (all that's required is to think that some such principle does exist), I briefly want to consider what strikes me as a promising candidate.

It is natural to think that in the context of legal trials, we should (unless we have reasons to believe otherwise) consider defendants as autonomous agents that are capable of freely determining their actions and overall conduct. In accordance with this assumption of autonomy, we should accept that agents are capable of diverging from the acts of their associates. As Wassermann (1991: 943) notes about the morality of statistical proof:

The basic idea is this: when we infer that the defendant acted like a majority of

⁶⁴ My emphasis.

the people in the stadium or prison yard, we treat him as someone randomly selected from the crowd, who can be assumed to have engaged in the modal behaviour... But [these inferences] are felt to be inconsistent with the law's commitment to treat the defendant as an autonomous individual, free to determine and alter his conduct at each moment.

Now, if we look at the situation in Prisoners, and more specifically Punish, we see that it undermines the principle of agent autonomy – the defendant's behaviour is inferred directly from the behaviour of the majority of prisoners and defendant's potential to diverge from the group's behaviour is ignored. The felt injustice of Punish, then, can be explained by the fact that it fails to treat the defendant as an autonomous individual - a *cornerstone* of our legal system. However, as mentioned before, while I take this to be a plausible explanation of why Punish strikes us as morally problematic, I am open to the idea that there are other equally plausible, or perhaps even more promising, ways of explaining the apparent injustice (or unfairness) of convicting individuals based on purely statistical evidence.

6.2. *General Upshots: The Legal vs The Epistemic*

Aside from providing proponents of Statistical Evidence with a response to the Conviction Argument, the previous considerations also point towards some more general conclusions about the connection between the legal and the epistemic domain. More specifically, the connection between the legal and the epistemic appears to be much less intimate than Littlejohn and others have suggested.

While many – including Lockean – think that our epistemic aim is to believe truly or to believe accurately, the law clearly has a *non-epistemic dimension* – i.e. it is not purely concerned with reaching true verdicts or with maximizing expected value. Amongst other things the law is concerned with reaching *just* and *fair* verdicts. Hence, our epistemic aim and the aim of the law can come apart. Sometimes, as cases like Prisoners illustrate, the epistemic aim of being an accurate epistemic agent conflict with our legal aims of having fair and just judiciary procedures. So, while epistemically speaking we may accept that our ultimate aim is truth, the aim of the

law is certainly not exhausted by truth-directed factors. Against this background, it shouldn't surprise us that our judgments about whether it is permissible to believe that a defendant is guilty can diverge from our judgments about whether it is permissible to convict a defendant – after all beliefs and legal convictions ultimately have different aims.

Once we acknowledge that our epistemic aim is different from our legal aims, we also have a more general reason to be suspicious about the overall strategy pursued by Littlejohn and others. The general thought underlying *The Conviction Argument* seems to be that we can somehow use legal considerations to evaluate epistemic principles and theories. In other words, opponents of Statistical Evidence are attempting to get epistemic mileage out of reflections on legal cases. However, once we see that our epistemic and our legal aims come apart, this strategy becomes less promising. Unless proponents of this strategy can convince us that the legal and the epistemic actually aim at the same thing, arguments against Statistical Evidence that try to draw epistemic conclusions from our judgments about legal cases are unlikely to succeed.

7. *The Argument from Comparative Probabilities*

Opponents of Statistical Evidence may insist that even if we accept *The Moral Argument against Punish* as a plausible response to what I called *The Conviction Argument*, there nevertheless remains a different kind of puzzle concerning legal cases, which proponents of the Lockean View cannot so easily escape. To set up the puzzle consider the following variant of the Prisoners.

Prisoners* 100 prisoners are exercising in the prison yard. Suddenly a prisoner attacks one of the guards. After the attack, one of the guards claims to have recognized the responsible prisoner – let's call him Bill – and offers to testify as an eyewitness in court. The environmental conditions on the day of the attack were just as they are on most other days.

Should Bill be convicted? I suppose many of us would judge that in this case it would be permissible to punish Bill. And again, our ordinary judgment is mirrored in our legal practices – the eyewitness testimony in Prisoners* would generally be regarded as sufficient for conviction. Likewise, most people would accept that in Prisoner* it would be permissible to believe that Bill was responsible for the attack. So far, so good.

However, opponents of Statistical Evidence may point out the following: of course eyewitness testimony is not perfectly reliable - perhaps the guard made a mistake - he may have been tired and mistaken Bill for someone else, perhaps there was another prisoner who looked just like Bill from the angle at which the guard was overlooking the prison yard, or maybe another prisoner decided to disguise himself to look like Bill so as to avoid being found guilty for attacking the guard. So, despite the fact that there is an eyewitness, we should assign the proposition that Bill attacked the guard a degree of probability less than 1. Let's say that a reasonable degree of probability to assign the proposition P*, that Bill was responsible for the attack, is 0.98.

But now opponents of Statistical Evidence may run the following argument. In Prisoner the probability that should be assigned to the proposition P, that the randomly chosen prisoner was involved in the attack, is 0.99. And in Prisoners* the probability we should assign to P* is 0.98. And, since $\Pr(P)=.99 > \Pr(P^*)=.98$, proponents of Lockean View seem committed to the idea that in Prisoners we are more justified in believing that the defendant is guilty than we are in Prisoners* - but this seems counterintuitive.⁶⁵ Surely, someone might argue, our belief that Bill in Prisoners* is guilty should be more justified than our belief that some randomly chosen defendant in Prisoners is guilty. Those who have this intuition can be seen as endorsing the following comparative claim.

The Comparative Claim We are more justified in believing P* than P.

⁶⁵ The precise degree of probability does not matter here. As we will see, for any degree of probability (< 1) one thinks we should assign to P*, we could alter the numbers of prisoners in Prisoners so that the reasonable degree of probabilities for P will exceed that of P*.

Since, this comparative claim seems to be in direct tension with the prediction of the Lockean View, its opponents may run the following argument.

The Comparative Probabilities Argument

- F1. If we accept the Lockean View, then we will be more justified in believing P than P*
- F2. However, we should accept The Comparative Claim– i.e. the claim that we are more justified in believing P* than P.
- Fc. Thus, we should reject the Lockean View.

How might Lockeans respond to this argument?

One option for proponents of the Lockean View would be to try and reject (F1), the claim that their view commits them to the result that in Prisoner we are more justified that the defendant is guilty than we are in Prisoners*. How might they do that?

First, they might point out that the Lockean View only tells us the conditions required for justification, it does not tell us anything about *degrees of justification* or how to rank two justified beliefs for their comparative strength. So, there is no direct route from the Lockean View to (F1). To get (F1), we require a further principle along the following lines:

Probabilistic Comparison For any two justified beliefs P and Q, if P's degree of probability exceeds Q's degree of probability, then P is more justified than Q.

Lockeans could resist (F1) of *The Comparative Probabilities Argument* by rejecting Probabilistic Comparison. But should they? Since Lockeans think that we should understand justification in probabilistic terms, I do not see on what grounds they may reject Probabilistic Comparison. Probabilistic Comparison seems highly plausible and if we take seriously the idea that justification amounts to high probability, then surely a higher degree of probability should amount to a higher degree of justification. Giving up on this idea seems to somehow betray the idea that justification is best

understood in probabilistic terms. Thus, I do not think that a rejection of Probabilistic Comparison is a promising option for Lockeans.

Without an obvious reason for rejecting (F1), it appears that the best option for Lockeans will be a rejection of (F2). However, making the claim that in Prisoners we are more are more justified in believing that the defendant is guilty than in Prisoners* will require some work. So, on what grounds might Lockeans argue that we should resist our initial inclination to accept The Comparative Claim?

Proponents of the Lockean View may point out that in Prisoners, P's degree of probability exceeds that of P* in Prisoners*. As such, the idea that we should be more justified in Prisoners than in Prisoners* seems sensible. The Comparative Claim suggests that we are more justified in believing a proposition that is less likely to be true – this should strike anyone who is committed to the idea that justified (or rational) belief aims at true as the wrong result. So, anyone who accepts that our epistemic aim is to believe truly – which Lockeans tend to do – will think that there are good reasons to reject our intuitions about the comparisons between Prisoners and Prisoners* expressed in The Comparative Claim; when it comes to comparisons between these two cases it seems that our intuitive judgments are at odds with the aim of believing truly. And since our intuitions are of course not infallible, it seems plausible that this is a case in which we have good reasons to reject them.

To make this response to *The Comparative Probabilities Argument* more compelling, proponents of the Lockean View will need to offer some explanation as to why people should have been so widely mistaken about their comparative judgments between cases involving statistical and non-statistical evidence like Prisoners and Prisoners*. In other words, Lockeans owe us some explanation as to why we should feel such a strong draw towards an allegedly bad principle – in this case The Comparative Claim. Fortunately, however, I think that there is a plausible story Lockeans can tell.

A plausible explanation for why we intuitively judge P* to be more justified than P is that we ordinarily don't think about beliefs based on perception (and presumably other types of evidence) in probabilistic terms. In the absence of any evidence to the contrary, it seems plausible that we assign beliefs based on perception

probability 1 – in other words we simply take them to be certain.⁶⁶ Why do we do this? Well, consider the circumstances under which our beliefs based on perception would be false; such circumstances will include unfortunate lighting, deception, the inference of hallucinatory drugs, etc. In our ordinary reasoning practices, we seem to simply *ignore* these possibilities or do not assign them any relevance. Another way of understanding this way of ignoring certain circumstances is to assign it probability 0. Thus, we judge that in light of the eyewitness testimony in Prisoners* the probability of P* is 1, while in Prisoners the degree of probability is 0.99. And we do this because we assign the interfering circumstances under which our belief that P* would be false - e.g. that the prison guard was tired and made a mistake identifying Bill, that the guard looked at the yard from a bad angle such that another prisoner looked just like Bill, or that another prisoner decided to disguise himself to look like Bill - probability 0.

Now, recognizing that this might explain our intuitive judgments in Prisoners and Prisoners* does of course nothing to justify this practice – all that the above considerations do is offer Lockeans with a plausible diagnostic to explain why we feel drawn towards The Comparative Claim, and therefore (F2) despite the fact that it recommends an – from a Lockean perspective - epistemically suboptimal position. Once we acknowledge that the guard in Prisoners* is fallible and recognize that the error-risk associated with P* may well be higher than the error-risk associated with P, we can come to appreciate the thought that in Prisoners we are more justified than in Prisoners* despite the fact that this goes against our initial intuitions. We might also come to appreciate why this reason against The Comparative Claim (and (F2)) may be more persuasive than the intuitions that support it.

Importantly then, while there may remain a small *residual problem* for Lockeans – after all they cannot accommodate the intuitive data captured in The Comparative Claim - it appears that they can offer compelling explanations for why this is the case and, perhaps more importantly, for why we should prefer their prediction over our intuitive judgment. Thus, this problem, we may call it the *Residual Problem*, should not be considered sufficient to undermine the Lockean View. In the

⁶⁶ A similar point has been made by Ross and Schroeder (2014).

end, Lockeans have a reasonable way of responding to *The Comparative Probabilities Argument* that is perfectly compatible with their commitment to the idea that justification amounts to something along the lines of high probability.

8. *So Far, it's Fine to be a Lockean*

In this chapter, we have introduced the Lockean View of justification and considered some of the leading arguments against it. The first three arguments have set out to undermine the Lockean View via Statistical Evidence. The idea was that if it can be shown that Statistical Evidence is unsustainable, then we can use it in a reduction on the Lockean View, which is committed to Statistical Evidence. However, as we have seen, none of the arguments against Statistical Evidence are particularly persuasive. Lockeans, and proponents of Statistical Evidence more generally, have more room for manoeuvring than their opponents seem to think. Next, we considered a more direct argument against the Lockean View – *The Argument from Comparative Probabilities*. This argument, while also involving evidence that is purely statistical in nature, purports to show that analysing justification in probabilistic terms clashed with our ordinary judgments about certain comparative cases. However, Lockeans again have a plausible way of responding to this argument. In the allegedly problematic cases, our intuitions recommend that we regard propositions that are less probable to be more justified than similar propositions that are more probable. Unsurprisingly, there is room for Lockeans to reasonably reject our intuitive judgments in these cases. Moreover, they have a plausible explanation for why our intuitions lead us astray in these cases. Once we appreciate what is going on in these comparative cases, the alleged argument against Lockeans loses much of its initial force. And even if we accept that it would be preferable if the Lockean View did not diverge from our ordinary judgments in these cases - which I don't think they should do – this would be nothing but a *small residual problem*. Importantly, none of the arguments considered in this chapter provide compelling reasons to reject the Lockean approach to justification. In the next chapter I consider a more pressing, and ultimately more successful, objection to the Lockean View.

Chapter 3

Multi Premise Closure and the Paradox of the Bitter Pill⁶⁷

1. *Multi Premise Closure*

In the previous chapter we introduced the most prominent picture of epistemic justification - the Lockean View - and considered but rejected one broad family of objections.

According to the very general formulation of the Lockean View I offered, a belief that P is justified for S iff P is highly probable for S.

Lockean View A belief that P is justified for S iff P is highly probable for S.

This probabilistic approach to justification does not just satisfy justification's core platitudes, but moreover, as Smith (2016: 29) recently pointed out, "there is, undeniably, something very natural about it."

In this chapter I consider a different and ultimately more successful objection to the Lockean View. As we will see, on probabilistic approaches to justification (with a threshold value $t < 1$), we cannot preserve a compelling deductive cogency constraint on justification. What does this mean?

It seems natural to suppose that if one is justified in believing some proposition, then one should also be justified in believing what logically (or deductively) follows from this belief. In other words, it seems plausible that if one is justified in believing a number of individual beliefs, then one should also be justified in believing what's logically entailed by these beliefs. The Lockean View however, is

⁶⁷ This chapter is based on my article "A bitter pill for closure," forthcoming in *Synthese*.

incompatible with this idea of deductive cogency. More precisely, the Lockean View is incompatible with the following compelling principle of multi premise closure (sometimes also called the Conjunction Rule⁶⁸ or the Conjunction Principle⁶⁹), according to which we are justified in believing the conjunction of our antecedently justified beliefs.

Multi Premise Closure (MPC) If S is justified in believing p and S is justified in believing q... and S is justified in believing n, then S is justified in believing the conjunction (p & q... n).

Before considering the arguments against the compatibility of the Lockean View with (MPC), it might be worth to briefly consider why being unable to preserve (MPC) is bad news for Lockeans. There are at least three reasons for thinking that rejecting (MPC) is problematic and that its preservation is a plausible *desideratum* for a satisfactory theory of justification.

First, (MPC) strikes many as intuitively compelling. Surely, if someone is justified in believing a number of individual propositions, then they should also be justified in believing these propositions together or at the same time. Nelkin (2000: 379) makes this point as follows, "... it is counterintuitive that one be rational as long as one is careful not to draw particular logical consequences from one's beliefs."⁷⁰ In short, one might think that a theory of justification is on the wrong track if it predicts that the property of *being justified* can be lost as the result of seemingly unproblematic logical operations like *conjunction introduction*.

A second, related, worry is that rejecting (MPC) has counterintuitive or odd *consequences* that one might reasonably want to avoid. Rejecting (MPC) makes room for the possibility that one can rationally accept, or justifiably believe, each premise of an argument and yet, without being irrational, deny the conclusion. Following an example by Smith (2016: 80), consider for instance the following exchange between

⁶⁸ See, for instance, Foley (1979).

⁶⁹ See, for instance, Evidine (1999) and Ryan (1991, 1996).

⁷⁰ A similar remark can be found in Ryan (1996: 124).

Tina, who is hosting a party to which she invited 100 guests, and Mike who is helping her organize it.

Mike: Do you think Tom is coming to the party?

Tina: Yes, he said he was.

Mike: Do you think Jess is coming to the party?

Tina: Yes, she RSVPed yesterday.

Mike: Do you think Ben is coming to the party?

Tina: Yes, he just phoned and confirmed.

... The conversation proceeds similarly for the remaining guests

Mike: Great, so you think everyone is coming?

Tina: No, I don't think that.

Mike: But you just said that you think Tom, Jess, Ben, and each of the other guests are coming.

Tina: That's right – I do think that each of them is coming but I do not think that all of them are coming.

This case illustrates the strange consequences of denying that justified belief is deductively closed - there seems to be something odd about Tina's response and the epistemic position that she is taking - it is difficult to see Tina's response as anything but irrational.⁷¹

Finally, it is sometimes argued that rejecting (MPC) amounts to a considerable theoretical cost. Why? (MPC) provides a natural method of expanding one's set of justified beliefs by generating new justified beliefs from old ones. But if we deny (MPC) then we lose this method for generating new justified beliefs. Put differently, by denying (MPC) we lose conjunction introduction as a general method of expanding our set of justified beliefs.

⁷¹ Someone might object that the dialogue does not explicitly use the word 'justified' and therefore fails to support the claim that rejecting (MPC) really has unattractive consequences. However, one can easily reinterpret the dialogue to include the notion of epistemic justification. The reason for the omission is that the dialogue is more natural without the explicit use of the word 'justification'.

To strengthen this argument, it might be helpful to remember the Permission platitude from chapter 1. According to Permission, if one is justified in believing a proposition P, then it is *at least* permissible for one to believe that P. Now, in light of this, (MPC) can be seen as a method for expanding one's set of permissible beliefs. This then allows us to understand the cost of denying (MPC) in terms of epistemic permissions: If we deny (MPC), then we lose a method for expanding the set of propositions we have permission to believe. In other words, the cost of denying (MPC) might become more obvious if it is understood as a loss of an epistemic permission-generating principle.

As we can see, there are number of reasons why we may want to preserve (MPC) and why we may consider it a strike against a theory of justification if it cannot preserve (MPC) - denying (MPC) does not just run against our intuitions but is also costly. Now, it is important to point out that (MPC) is of course not a core platitude of justification – if a theory of justification cannot preserve (MPC), then surely we shouldn't conclude that it therefore fails to be a theory of justification.⁷² Instead, we should include (MPC) amongst the list of plausible desiderata for a satisfactory theory of justification. In other words, we should think of (MPC) as a feature that we should try to preserve if possible. Similarly, it should count in favour of a theory if it is able to preserve (MPC).

2. *The Ingredients for Closure Paradoxes*

The claim against the compatibility of Lockean Views with (MPC) is often made via two epistemic paradoxes - the *lottery paradox*⁷³ and the *preface paradox*⁷⁴. These paradoxes rely on three epistemic principles. The first one, (MPC), we already introduced. The other two principles are the following.

⁷² As mentioned in chapter 1, Leplin (2009) denies this sort of claim. He argues that justification is an *adequacy condition* for theories justification. However, once we recognize the contention and the pervasive disagreement over the status of (MPC) this claim starts to look unacceptable.

⁷³ The lottery paradox was first introduced by Kyburg (1961).

⁷⁴ The lottery paradox was first introduced by Makinson (1965).

Sufficiency Thesis (ST) If p is highly probable for S , then S is justified in believing p .⁷⁵

No Contradiction (NC) S is never justified in believing $(p \text{ and } \sim p)$.

(ST) is simply the sufficiency direction of the Lockean View. Since the necessity direction is not required to get the paradoxes off the ground it seems appropriate, in the interest of simplicity, to present the sufficiency thesis as a separate principle.⁷⁶ (NC) states that one cannot justifiably believe a contradiction of the form $(p \text{ and } \sim p)$. (NC) is not just intuitively compelling but also widely regarded as one of the most fundamental constraints on rationality and epistemic justification. As a result, it is generally considered to be beyond reproach and for the purpose of this study I will follow this orthodoxy and regard (NC) as non-negotiable.

What the lottery and preface paradox set out to show, is that the three principles - (MPC), (ST), and (NC) – while individually plausible are jointly inconsistent. In other words, one of them needs to go. For Lockeans this seems to suggest that they need to reject (MPC). However, recognizing that this would be a cost for their view, a number of strategies have been proposed that offer the promise of solving the two paradoxes without denying (MPC). As a result, the lottery paradox and the preface paradox have been unable to make a decisive case for the claim that Lockeans cannot preserve (MPC). The primary objective of this chapter is to introduce a new epistemic paradox that makes a stronger, and ultimately more decisive, case against the compatibility of Lockean approaches to justification with (MPC). In the next two sections I introduce the lottery and the preface paradox and outline the two

⁷⁵ The term ‘Sufficiency Thesis’ has previously been used by Douven&Williamson (2006) and Kelp (2017). Sometimes the lottery and preface paradox are presented in terms of rational belief, rather than justified belief. In these cases the paradoxes is set up using the Lockean Thesis, which explains rational belief in terms of rational degrees of confidence or credences.

Lockean Thesis It is rational for you to believe p just in case it is rational for you to have degree of confidence y in p , where $y > x$. (Foley 1992: 112)

Since (ST) is essentially just a justification analogue of the Lockean Thesis, I will, for the purposes of this study, not distinguish between these two notions.

⁷⁶ Especially, as we will see later, one might want to try and reject only the sufficiency direction of the Lockean View but keep the necessity direction.

leading solutions available to those wishing to preserve (MPC). I then introduce a new, structurally related, paradox that is immune to these closure-preserving solutions. I will call this paradox, *The Paradox of the Pill*. In addition to presenting a more robust counterexample to (MPC), the new paradox also reveals that the strategies that were previously thought to get closure out of trouble are not sufficiently general to achieve this task, as they fail to apply to similar closure-threatening paradoxes in the same vicinity.

3. *The Lottery Paradox*

Consider Juliet who enters a fair lottery with n tickets and a guaranteed winner. Assuming that Juliet knows these facts about the lottery, then for any threshold value $t < 1$ used to spell out ‘high probability’ in (ST) we can always construct a lottery large enough to create the following paradox.

- (L1) For any ticket Juliet is justified in believing that the ticket will lose; $JB_J(t_1) \& JB_J(t_2) \dots \& JB_J(t_n)$.
- (L2) If Juliet is justified in believing of any given ticket that it will lose, then Juliet is justified in believing that all tickets will lose; $JB_J(t_1 \& t_2 \dots \& t_n)$.
- (L3) Since Juliet knows that the lottery is fair and that it has a guaranteed winner, Juliet is justified in believing that at least one ticket will win; $JB_J \sim (t_1 \& t_2 \dots \& t_n)$.

When combined with (MPC), (L1) through (L3) entail a violation of (NC) as they give rise to the following contradiction.

- (L4) Juliet is justified in believing that every ticket will lose and that one ticket will win; $JB_J((t_1 \& t_2 \dots \& t_n) \& \sim (t_1 \& t_2 \dots \& t_n))$.

Since accepting the contradiction in (L4) is unacceptable as it would mean rejecting (NC) and rejecting (L3) is not a live option given the setup, we are left with two alternatives: either we reject (L1), the claim that Juliet is justified in believing of any

ticket in the lottery that it will lose, or we reject (L2), the claim that Juliet is justified in believing that all tickets are losers. The former requires rejecting (ST)⁷⁷, while the latter requires rejecting (MPC).⁷⁸ Thus, if (MPC) is to be preserved Lockean will need to deny (L1) and insist that lottery beliefs, no matter how probable, fail to be justified.

Denying justification for lottery propositions is of course at odds with (ST) - after all lottery propositions can be highly probable - and proponents of this strategy must somehow reconcile this tension. This is usually done in one of two ways: (a) by proposing modifications to (ST) or (b) by rejecting (ST) altogether and replacing it with an entirely different notion of justification. Whilst Leplin (2009) and Smith (2010, 2016) recently opted for a more radical strategy along the lines of (b), this chapter is primarily concerned with the more common and less radical *modificationist* strategies of type (a).⁷⁹ The result of these modifications is usually to prevent beliefs based on purely statistical evidence from counting as justified. Following this strategy, we get a modified version of the Sufficiency Thesis along the following lines.

Sufficiency Thesis* (ST*) If p is highly probable for S and further conditions are met, then S is justified in believing p.

One influential closure-preserving strategy along these lines is due to Dana Nelkin (2000).⁸⁰ Nelkin (388) argues that what initially motivates an acceptance of (L1) is a certain inference pattern, which she calls p-inferences.

(P-inference) p has a statistical probability of n [where n is a very high number] → p

⁷⁷ Proponents of this strategy include Ryan (1996), Kaplan (1996), Nelkin (2000), Leplin (2009), Buchak (2014), Staffel (2015), and Smith (2010, 2016)

⁷⁸ Proponents of this strategy include Kyburg (1961), Foley (1992, 2009), Christensen (2004), and Sturgeon (2008)

⁷⁹ We will consider strategies of type (b) in chapter 4.

⁸⁰ Similar views are defended or echoed in Kaplan (1996), Enoch et al. (2012), Buchak (2014), and Staffel (2015).

This inference pattern however, Nelkin argues, should be rejected. In order for a belief to be justified, so the argument continues, it requires more than just statistical probability. More specifically, in order for a belief to be justified according to Nelkin an agent must also be in a position to *suppose* that there exists a causal or explanatory connection between the belief and the facts that would make it true. Consider for instance beliefs based on perceptual experiences. For such beliefs, we are usually in a position to suppose that there exists a causal connection between the belief and the facts that would make the belief true, i.e. the objects perceived. For beliefs about lottery tickets on the other hand this condition is not met, as there is no causal connection between the belief that your ticket will lose and the fact that would make the lottery belief true, i.e. the drawing of the lottery; We would believe that our ticket was a loser even if it was the case that we held the winning ticket. As a result, Nelkin argues, beliefs in lottery propositions are not justified. Thus, by modifying the Lockean Thesis to rule out beliefs based on statistical evidence, the lottery paradox can be solved whilst closure is preserved – let’s call solutions to the lottery paradox following this strategy *Solutions from Statistical Evidence*

This strategy for responding to the lottery paradox can easily be turned into another argument against justification from statistical evidence discussed in the previous chapter. We can easily imagine someone arguing as follows: (1) If we reject that beliefs based on purely statistical evidence can be justified, then we can preserve (MPC). (2) We should preserve (MPC). (3) Therefore, we should reject that beliefs based on statistical evidence can be justified. While this argument, as we will see later, ultimately fails, let’s accept it for the time being.⁸¹ Lockeans who want to avoid the consequence of having to deny (MPC) might take this argument as motivation to modify their view in the way outlined above. If it turned out that denying statistical evidence would put Lockeans in a position to preserve (MPC), so the thought goes,

⁸¹ The argument ultimately fails because (1) is false. The reason for this will become apparent in the upcoming discussion.

then it would arguably be worth it to make slight modifications to (ST) – and subsequently the Lockean View - to deliver this result.⁸²

Before moving on the preface paradox, it is important to draw our attention to one final, yet crucially important, observation. Note that solutions that set out to solve the lottery paradox by rejecting (L1) are motivated by, or at least rely on, the fact that it strikes us as intuitively *acceptable* to deny justification for lottery beliefs. This is an important feature underwriting these solutions. Douven (2008) recently made the stronger claim that denying justification for lottery beliefs does not just strike us as an intuitively *acceptable* response, but that it strikes many as the intuitively *correct* response; “Many have the intuition that the right response to the Lottery Paradox is to deny that one can justifiably believe of even a single lottery ticket that it will lose” (204).⁸³ Now, we know from the previous chapter that this stronger claim is problematic for a number of reasons. First, while there may be some people who find this solution intuitive, there is also a considerable number of people who find this solution counterintuitive. And second, a number of empirical studies have consistently shown that people ordinarily do think that we can justifiably believe of individual lottery tickets that they are going to lose. However, while I do not think that the stronger claim that rejecting (L1) strikes many as the correct solution is very compelling, I do think that its weaker counterpart, namely that it strikes many as an intuitively *acceptable* solution, is correct - especially if this solution will allow us to preserve (MPC). Why is this an important feature? If it was the case that denying the justificatory status of lottery beliefs struck many as unacceptable, then solutions the lottery paradox that deny (L1) would appear unsatisfactory (or at least unpersuasive).

4. *The Preface Paradox*

⁸² One might object that at this stage Lockeans would betray the main tenet of their view, namely that justification is best explained in terms of high probability. As such, this modificationist strategy runs the risk of looking somewhat unprincipled. Here is how Lockeans who want to preserve (MPC) might respond: The benefit of preserving (MPC) far outweighs the worries that the means of securing this result might run the risk of appearing slightly unprincipled. However, since this strategy will in the end be unpromising anyways, we should not be too worried about this issue.

⁸³ It is important to point out that Douven is only summarizing the state of the debate; he is not endorsing this as the correct solution to the lottery paradox.

The preface paradox is structurally related to the lottery paradox as it too involves a combination of (ST), (MPC), and (NC) to create a seemingly inconsistent set of beliefs. The paradox is often presented as follows.

An author has just finished an ambitious, yet carefully researched, book. On the basis of her research the author has good evidence for every claim made in the book and is justified in believing that each claim made in the book is true. However, the author is also aware that when it comes to ambitious books, even the best scholarship can turn out to contain errors and in the past this has always been the case. In light of this the author includes in the preface of her work an apology for any errors in her book.

The author's epistemic position is considered to be problematic as she appears to be justified in believing that each claim in her book is true whilst also being justified in believing that there is at least one error in her book. The formal structure of the paradox can be captured as follows.

- (P1) The author is justified in believing each claim made in the book; $JB_A(p_1) \& JB_A(p_2) \dots \& JB_A(p_n)$.
- (P2) If the author is justified in believing each claim made in her book is true, then the author is justified in believing that all claims made in her book are true; $JB_A(p_1 \& p_2 \dots \& p_n)$.
- (P3) The author is justified in believing that her book contains at least one error; $JB_A \sim (p_1 \& p_2 \dots \& p_n)$.

However, taken together (MPC) and (P1)-(P3) entail a violation of (NC) as they combine to produce the following contradiction.

- (P4) The author is justified in believing a contradiction, namely that all claims in her book are true and that her book contains at least one error; $JB_J((p_1 \& p_2 \dots \& p_n) \& \sim (p_1 \& p_2 \dots \& p_n))$.

As in the lottery paradox, there are two prominent solutions to the preface paradox. One option is to reject (P2), which would effectively deny (MPC). Alternatively, those sympathetic to closure have the option of rejecting (P3) instead. Advocates of this latter strategy include Pollock (1986), Ryan (1991), Leplin (2009), Kaplan (2013), Kim (2015), and Smith (2016).

Arguments for rejecting (P3) typically pivot on the idea that the author's mere recognition of her own fallibility is not sufficient to justify the belief that her book will contain any errors. Given that our author has good reasons supporting every claim in the book, mere doubt arising from recognizing her fallibility is insufficient to defeat justification for believing the conjunction. To support this claim, Kaplan emphasizes the lack of a *special reason* for believing that the book contains any errors, "... [the author] should stand by her book – she should not shrink from saying that everything in it is true – until she has found *special reason* to think it contains an error" (2013: 16).⁸⁴

Importantly, according to Kaplan and others following this strategy, the author is of course perfectly justified in believing that her book will likely contain errors; "[The author] is even free to confess to *being confident* that her book will turn out to contain some error or other. The only thing she cannot do is flatly take back what she's said in the rest of the book: she cannot flatly say (or say anything that entails) that her book *has errors*." (2013: 17)⁸⁵ Similar remarks can be found in recent work by Kim, "... while authors are reasonable to assert that *surely* there are errors, the reasonableness of this assertion does not imply that they are rational to believe that there *are* errors" (2015: 1024).⁸⁶ Thus Kaplan and Kim invite us to replace (P3) with a weaker premise along the following lines.

⁸⁴ Emphasis in italics added.

⁸⁵ Emphasis in italics added.

⁸⁶ Emphasis in italics added.

(P3*) The author is justified in believing that it is likely that errors will be found in her book.

With (P3*) in place the paradox is resolved because there is no contradiction in our author believing that all claims made in her book are true whilst also believing it to be *likely* that her book contains at least one error.

Thus, it seems that proponents of closure can once again defuse the challenge of a supposedly closure-threatening paradox. Proponents of (MPC) can exploit the fact that in the preface paradox there is no *special reason* for believing that the conjunction in question is false. Let us call solutions of this ilk, *Solutions From Defending the Conjunction*.

It is important to point out that this strategy would again require modifications to (ST) and thereby the Lockean View.⁸⁷ Why? According to (ST) it is precisely the property of being highly probable for S that makes a belief justified. Since the belief that the book contains at least one error is highly probable (ST) would predict that it is justified. In other words, given how (ST) is formulated it is not clear how (P3*) could be accommodated on the standard Lockean Picture. There is however a nice way of unifying the strategy proposed above with the *Solution from Statistical Evidence* previously considered in response to the lottery paradox. According to *Solution from Statistical Evidence* we need to modify (ST) so that it rules out justification for beliefs based on purely statistical evidence. Similarly, according to *Solutions From Defending the Conjunction*, justification requires a special reason in support of the proposition. Now, presumably for beliefs based on purely statistical evidence, one also lacks a special reason for believing that it is true. After all there is no *special reason* for thinking that one's lottery ticket is going to lose; all one has is general evidence that applies equally to all other lottery tickets. So, depending on how proponents of *Solutions From Defending the Conjunction* spell out the *special reason* requirement, it is likely to also rule out justification for lottery beliefs. In other words, *Solutions From Defending the Conjunction*, like *Solutions From Defending the*

⁸⁷ At this point one could raise a similar worry to the one expressed in footnote 82, namely that the required modifications will betray the Lockean's commitment to the idea that justification is best analyzed in terms of high probability. However, since we will later reject these strategies anyways in favor of rejecting (MPC), we do not need to worry about this issue here.

Conjunction, will require modifications to (ST) along the lines of (ST*). The hope then is Lockeans who want to preserve (MPC) may find one candidate for the ‘further condition’ slot of (ST*) that will be able to solve both the lottery and the preface paradox.

5. *The Possibility of a Third Related Paradox*

Thus far I have argued that neither the lottery paradox nor the paradox of the preface are able to make a decisive case against the claim that justification is closed under multi premise deduction. As we have seen, each paradox has a certain feature that Lockeans sympathetic to (MPC) can reasonably exploit in responding to the respective puzzles. In the lottery paradox proponents of closure can exploit the fact that the evidence in question is purely statistical and resort to *Solutions from Statistical Evidence*. Likewise, in the preface case proponents can appeal to *Solutions From Defending the Conjunction* seeing that there is no special reason for believing that the conjunction is false.

At this point one might wonder if it is possible to construct a new closure threatening paradox that would be immune to *Solution from Statistical Evidence* as well as *Solutions From Defending the Conjunction*. Such a paradox would have to satisfy the following two conditions: (a) the beliefs in the individual conjuncts must not be based on purely statistical evidence and (b) there must be a special reason for believing that a sufficiently long conjunction will contain an error. We can expect a paradox with these features to present a more robust counterexample to the claim that justification is closed under deduction than its two predecessors. However, currently there is no paradox available that has precisely this mix of features.⁸⁸ The next section will present such a paradox and fill this lacuna.

⁸⁸An anonymous referee for *Synthese* helpfully suggested that a certain variant of the preface paradox - the Homogeneous Preface Paradox - recently proposed by Easwaran and Fitelson (2015) might meet these two conditions. In what follows I will briefly argue that there are very good reasons for thinking that the Homogeneous Preface Paradox, like the standard preface paradox, fails to meet condition (b). As a result, proponents of (MPC) will be able to respond to the Homogeneous Preface Paradox by appeal to *Solutions From Defending the Conjunction*. In the Homogeneous Preface Paradox, unlike in the standard preface paradox, (P3) is not just supported by general second-order evidence about our human fallibility; instead the author, John, an empirical scientist, writes an ambitious book about the very hypothesis (H): that all scientific/empirical books of sufficient complexity contain at least one

6. *A Bitter Pill for Closure*

Consider the following case.

S is given a bitter pill that ensures that a very small portion of S's ordinarily justified beliefs, let's say 1 out of every 10,000, chosen at random, will be false. The pill achieves this result by occasionally impairing S's cognitive connection to the evidence resulting in occurrences of, for instance, misperceptions or false memories. Importantly however these occurrences are incredibly rare. Finally, S knows about the effects of the bitter pill.⁸⁹

Let's briefly consider S's epistemic situation after taking the pill. After taking the pill S continues to form beliefs like everyone else and the vast majority of those beliefs, will turn out to be true - S will continue to form beliefs based on perception, memory, and testimony, etc. just as before taking the pill. However, not all of S's justified beliefs will be true, for the pill ensures that a small number of beliefs, chosen at random, are going to be false. Considering that these occurrences are extremely rare,

false claim (2015: 10). So, John, rather than having only second-order evidence in support of the claim that his book contains at least one error, now infers the belief from the first-order claims in his book. Does this change in the *type* of evidence supporting (P3) amount to there being any *special reason* for thinking that *John's book* contains any errors? Arguably not. Sure, John has written a book rich in first-order claims supporting (H), a hypothesis about the general fallibility of ambitious scientific works, but these claims do not provide any *special reason* in support of the claim that *John's book* contains any error. As Kaplan suggests, "even if there are other authors of comparably ambitious works, no less careful (and, perhaps, some more meticulous) than [the author], who have nonetheless failed to write error-free books. Those aren't special reasons for not standing behind [the author's] book; they are reasons that are routinely present..." (2013: 16). In the case of John then, the first-order claims supporting (H) fail to be special reasons for believing that his book contains any errors, for the claims supporting (H) will be routinely present as they bear no special relation to John's current book - they will apply equally to his next book. Hence, there are strong reasons for thinking that the Homogeneous Preface Paradox does not meet condition (b). Moreover, in the absence of any special reason for believing that the book contains at least one error, proponents of (MPC) can respond to the Homogeneous Preface Paradox by appealing to *Solutions From Defending the Conjunction*.

⁸⁹ As with most thought experiments it is of course perfectly possible to provide an alternative narrative to capture the same epistemically relevant features. Those who find the pill case conceptually problematic are free to come up with an alternative narrative as long as it incorporates the same epistemic features – one obvious candidate would be a malevolent demon, which decides to deceive an agent about a small number of randomly chosen beliefs.

and that the vast majority of S's beliefs will remain entirely unaffected, we should expect that the overall impact of the pill on S's epistemic situation should be rather minimal. After all, it is likely that we already form and sustain a host of false beliefs anyways and that we are no strangers to the occasional adaptation of a false belief. Moreover, recall that in the last chapter we established that fallibility is a widely accepted constraint on theories of justification. Hence, (ST) will plausibly predict that those of S's beliefs that are made sufficiently probable by S's evidence will continue to be justified.

From the pill case we can derive a puzzle similar to the lottery and the preface paradox. Let's call this new paradox, *The Paradox of the Pill*.

(BP1) S, after taking the pill, continues to be justified in believing the propositions made sufficiently probable by her evidence; $JB_S(p_1) \& JB_S(p_2) \dots \& JB_S(p_n)$.

(BP2) If S continues to form justified beliefs, then S is justified in believing the conjunction of her individually justified beliefs; $JB_S(p_1 \& p_2 \dots \& p_n)$.

(BP3) S is justified in believing that the conjunction of her justified beliefs will contain at least some false belief; $JB_S \sim (p_1 \& p_2 \dots \& p_n)$.

The combination of (BP2) and (BP3) however violates (NC) as it commits S to an inconsistent set of beliefs. Under closure S is justified in believing the conjunction of her individually justified beliefs, whilst the effects of the pill also justify her in believing that a long enough conjunction of her individually justified beliefs is going to be false.

(BP4) S is justified in believing a contradiction, namely that the conjunction of her justified beliefs is true and that some members of the conjunction are false; $JB_S((p_1 \& p_2 \dots \& p_n) \& \sim (p_1 \& p_2 \dots \& p_n))$.

As in the previous cases, solving the paradox will require rejecting either (BP1), (BP2), or (BP3). If (MPC) is to be preserved, then (BP2) must be retained. This leaves

(BP1) or (BP3). However, neither of the popular solutions for solving closure puzzles will deliver this result; neither *Solutions from Statistical Evidence*, which would target the individual conjuncts in (BP1), nor *Solutions From Defending the Conjunction*, which would target (BP3), will apply to the new paradox.

Solutions from Defending the Conjunction, which were plausible in response to the preface paradox, turn out to be unavailable in the case of the pill paradox. Recall, that in response to the preface paradox Kaplan and others argued that there was no special reason for believing that the conjunction would contain any errors. It was this feature of the preface paradox that proponents of (MPC) exploited to reject (P3). In the paradox of the pill however, there does exist such a special reason - the effects of the pill guarantee that a sufficiently long conjunction will contain at least one false belief. As a result, *Solutions from Defending the Conjunction* do not provide the necessary resources for rejecting (BP3) in the new puzzle.

Similar worries arise for *Solutions from Statistical Evidence*. Recall that *Solutions from Statistical Evidence* attempt to solve the lottery paradox by denying that purely statistical evidence, no matter how strong, is sufficient for justification. An influential argument for this type of strategy was due to Nelkin who argued that in order for a belief to be justified an agent must be in a position to *assume* that there exists a causal or explanatory connection between the belief and the evidence that will make it true. In case of the paradox of the pill however these types of solutions will fail. In the paradox of the pill, S's beliefs are not based on purely statistical evidence, as S continues to form beliefs based on perception, testimony, memory, etc. And since S's beliefs are formed using these ordinary methods, for every individual belief, S will be in a position to *suppose* that there exists a causal or explanatory connection between her belief and the facts that would make the belief true. Furthermore, consider that even under ordinary circumstances our beliefs are liable to occasional errors. We would not however conclude that this gets in the way of supposing that generally there exists a causal connection between our beliefs and the facts that make it true. Since we are no strangers to the occasional formation of false beliefs and since this does not compromise the assumption that our beliefs are causally connected to the evidence, there is no reason to deny that after taking the pill the agent can assume, for each belief, that there exists a causal connection between the belief and the facts that would

make it true. Thus, *Solutions from Statistical Evidence* will also not be able to solve the new paradox.

Since the prominent closure-preserving solutions fail to apply to the new puzzle, it appears that the most promising way out of the paradox is to reject (BP2), and as a result (MPC). To cement this conclusion let's briefly consider the prospects of alternative strategies for rejecting (BP1), i.e. strategies that would solve the lottery paradox by rejecting (L1) and that would in principle apply to the new paradox. Pollock (1983) and Ryan (1996), for instance, reject (L1) on the more general ground that whenever one has a set of beliefs and it is known that the set contains at least one false member, then, assuming one does not know which of the members is false, one is not justified in believing *any* member of the set.⁹⁰ In the case of the new paradox, these considerations might motivate a rejection of (BP1) by insisting that after taking the pill *none* of S's beliefs are justified. However, recall that strategies akin to *Solutions from Statistical Evidence* that reject (L1) by denying that we can have justified beliefs about lottery tickets are motivated by, or at least rely on, the idea that denying justification for lottery beliefs strikes us as an intuitively acceptable response. Douven made the stronger claim that denying justification for lottery beliefs is not just an intuitively acceptable response, but that it is the *intuitively correct* response. In case of the Paradox of the Pill however, insisting that *all* of S's beliefs are unjustified no longer strikes us as an intuitively correct, or even intuitively acceptable, result. Seeing that the pill only affects 0.01% of beliefs, it appears that embracing global skepticism about justification (and presumably knowledge) is an excessively costly and unconvincing response to the new puzzle; especially if we recall that most accounts of justification depart from the idea that justification is fallible, i.e. that we can have justified false beliefs. Thus, even strategies that solve the lottery by rejecting (L1) and that would in principle apply to the new paradox by rejecting (BP1) will not get closure out of trouble. In the case of the new paradox, these solutions lose their intuitive plausibility, which is what made them promising closure-preserving solutions to the Lottery Paradox in the first place.

⁹⁰ To motivate this claim, Ryan (1996) appeals to a principle she calls the 'Avoid Falsity Principle' while Pollock (1983) introduces the idea that became known as 'Collective Defeat'.

A final response to the Paradox of the Pill might try to avoid the skeptical implication of the above strategies by proposing that only *some* of S's post-pill beliefs will be unjustified.⁹¹ This, one might hope, would provide a more plausible response to the new paradox. However, it is not immediately obvious how this solution may be motivated. After all, the effects of the pill are both *global* so that the pill could affect any of S's beliefs, and *random*. This means that for every belief, regardless of its precise degree of evidential support or how confident S is about its truth, there exists an equal chance that it is affected by the pill. Put differently, we might say that with regards to the pill, all of S's beliefs are epistemically on par. As a result, it is not immediately obvious how to make sense of the idea that only some of S's post-pill beliefs are unjustified. However, Kroedel (2012) recently proposed a solution to the lottery paradox that promises to provide the resources for such a response. But, as we will see, this solution, even if we accept it as a possible solution to the lottery paradox, seems less promising in response to the new Paradox of the Pill.

Kroedel suggests that if we suppose that epistemic justification is a species of permissibility, then we can solve the lottery paradox without giving up (MPC).⁹² On a permissibility reading of justification, (L2) turns into the following claim:

(LP2) If Juliet is permitted to believe that t_1 will lose, then Juliet is *permitted* to believe of any ticket in the lottery that it will lose.

This claim however, Kroedel argues remains ambiguous. Let PE be a permissibility operator, such that $PE_j Bp$ denotes the sentence 'It is permissible for Juliet to believe p'. Kroedel argues that depending on the scope of PE (LP2) can have a narrow and a wide-scope interpretation.

The narrow interpretation only permits *separate* beliefs in lottery propositions.

⁹¹ I want to thank two anonymous referees for Syntheses for bringing this proposal to my attention.

⁹² Recall from chapter 1 that there are some general concerns about the plausibility of the view that epistemic justification can be interpreted as *mere* permissibility claims. White (2005), for instance, explicitly rejects this idea and argues that justification amounts to a deontologically stronger 'ought' claim. For our purposes here, I will simply grant that considering epistemic justification a species of permissibility is a plausible strategy.

(Narrow) $PE_J Bp_1 \ \& \ PE_J Bp_2 \dots \ \& \ PE_J Bp_n$

Under the Wide-scope reading of (LP2) permissions agglomerate and one is permitted to believe all the propositions together.

(Wide) $PE_J(Bp_1 \ \& \ Bp_2 \ \dots \ Bp_n)$

With this distinction in place we have two possible interpretations for (LP2).

(LP2*) If Juliet is permitted to believe that t_1 will lose, then Juliet is permitted to believe of any ticket in the lottery that it will lose; $PE_J Bp_1 \ \& \ PE_J Bp_2 \dots \ \& \ PE_J Bp_n$.

(LP2**) If Juliet is permitted to believe that t_1 will lose, then Juliet is permitted to believe of any ticket in the lottery that it will lose; $PE_J(Bp_1 \ \& \ Bp_2 \ \dots \ Bp_n)$.

The lottery paradox only arises if we accept the wide reading, i.e (LP2**). Under the wide reading we get the paradoxical conclusion that Juliet is permitted to believe that every ticket will lose and also permitted to believe that one ticket will not lose. In this respect (LP2**) is relevantly similar to (L2) from the initial discussion of the lottery paradox above. The question that naturally arises at this point is if there are any good reasons to reject the wide-scope reading and to adopt the narrow-scope reading instead? If there is a good reason to prefer the narrow-scope to the wide-scope reading, then we have a way out of the lottery paradox.

Kroedel argues that looking at how permissions work, provides us with a good reason for favouring the narrow-scope reading over the wide-scope reading – i.e. it provides a reason for favouring (LP2*) over (LP2**). Here is why. One important feature about permissions is that they *do not* agglomerate. Even though Juliet has permission to believe of any given ticket that it will lose, she does not therefore have permission to believe that all tickets are going to lose. Permissions simply do not work this way. Kroedel uses a helpful analogy to illustrate this point. “For instance, I might be permitted to eat this piece of the cake, permitted to eat that piece of cake, etc. without being permitted to eat the whole cake” (2012: 59). Given that permissions do

not agglomerate, Kroedel argues, we should reject the wide scope reading (LP2**) and instead accept the narrow scope reading (LP2*). With (LP2*) in place the lottery paradox is resolved because Juliet lacks permission to believe that every ticket is a losing one.

At this point one might wonder whether Kroedel's account does not simply amount to a rejection of (MPC). After all Juliet has permission to believe of every ticket that it will lose but lacks permission to believe the conjunction of these claims. This might initially strike us as a rejection of closure. Kroedel's account however does not need to reject (MPC). The reason one does not have permission to believe the conjunction of individually permitted beliefs is explained by the fact that permissions do not agglomerate. This is simply a deontic feature of permissions and therefore *independent* of concerns about closure. Kroedel makes this point explicit, "Since each lottery ticket has a high probability of losing, I am permitted to believe that Ticket #1 will lose, permitted to believe that that Ticket #2 will lose, etc. But since permissibility does not agglomerate, it does not follow that I am permitted to have all these beliefs at once. Neither does it follow that I am permitted to have a single belief whose content is the conjunction of the contents of these beliefs- even if we uphold the closure principle" (2012, p. 60). In short, the reason Juliet is not justified in believing the conjunction of her individual is simply a feature of how permissions work, regardless of what we make of (MPC).

Let us consider how Kroedel's permissibility strategy would apply to the Paradox of the Pill. While S continues to form justified beliefs after taking the pill, on a permissibility reading all that this amounts to, is the notion that S continues to acquire permissions to believe certain propositions. However, permissions do not agglomerate and therefore S is not permitted to believe that all her beliefs are true. This seems to be the correct result seeing that S knows that 0.01% of her beliefs are guaranteed to be false. Thus, on a permissibility reading of justification we can deny (BP2), as S is not permitted to believe that all of her beliefs are true. The paradox seems to be avoided.

There is however a serious obstacle for this strategy. Given that permissions do not agglomerate and that one does not have permission to believe all one's individually permitted beliefs together, one will have to somehow identify some post-

pill beliefs about which to suspend judgment. More precisely, one will have to identify some of the proposition that one has permission to believe and refuse to believe them. But how do we go about picking these beliefs? In lottery cases this is an easy task - Kroedel suggests that one can just *arbitrarily* chose a ticket at random and suspend judgment about whether or not it is a loser.

Consider a lottery with only three tickets. It is permissible for me to believe of any of the three tickets that it will lose. However, knowing that permissions do not agglomerate I *arbitrarily* chose to believe that ticket #3 will lose and suspend judgment about the other two. Kroedel explicitly endorses this type of arbitrariness, “What is arbitrary is that I come to believe of this specific ticket that it will lose. It is hard to see what is supposed to be bad about this sort of arbitrariness, ...” (2012, 60). This arbitrariness in picking and choosing which of the individually permissible beliefs to adopt and which ones to suspend judgment about might work in lottery cases where there is a strict homogeneity between the individual proposition. However, it is not clear that this approach will be acceptable in other cases. More specifically, this strategy seems to become problematic in cases where the evidence is non-statistical.

Eder (2015) previously argued that precisely this arbitrariness of picking and choosing what to believe and what not do believe prevents Kroedel’s solution from applying to the preface paradox in which the evidence is not statistical. In the preface case Eder argues, it would not be a palatable option to suppose an author should arbitrarily identify one claim in the book and simply suspend judgment about it. After all, in the preface paradox the author has good evidence supporting each claim in the book; as Eder (2015) warns, “Ignoring this and choosing arbitrarily which propositions not to believe seems to make dealing with rational belief (evidential support, and rational degrees of belief) somewhat superfluous: this is not a palatable option here...” (669). In short, there would be something epistemically disingenuous about an author who is adopting this stance towards their book. It is easy to see how a similar worry about the generality of Kroedel’s solution extends to the case of the pill paradox.

Consider S, after taking the pill, and the vast amount of permissions that she has to believe various propositions. Given that permissions do not agglomerate S knows that she cannot rationally believe all those propositions together. Thus S, in an

effort to be epistemically rational, will need to choose some propositions that she, despite being permitted to believe them, will suspend judgment about. But how should S go about identifying these beliefs in the pill scenario? It appears that we will struggle to identify, in a principled manner, which propositions S should not believe. Kroedel's suggestions of simply choosing some propositions at random does not seem appropriate in this case. Overall then, Kroedel's account fails to provide a satisfactory solution to the bitter pill paradox.

These considerations point towards the following strategy for dealing with the Paradox of the Pill: either we grant that S has justification for every belief that is probable on her evidence, or, if one thinks that the right response to the paradox somehow involves rejecting (BP1), then one should be prepared to deny justification for all of S's beliefs. Thus, aiming to solve the new paradox by targeting (BP1), i.e. by seeking fault with the individual conjuncts, is likely to be unsatisfactory. As pointed out before however, many of the popular strategies for solving the lottery paradox by targeting (L1) do not apply to (BP1) in the first place.

Seeing that many of the prominent closure-preserving strategies are not equipped to deal with the new puzzle, we can conclude that The Paradox of the Pill makes a significantly stronger case against (MPC) than its two predecessors.

7. *Generality Worries for Standard Closure-Preserving Solutions*

It appears that *The Paradox of the Pill* is not just bad news for closure. The new puzzle also licenses some more general concerns about certain strategies for dealing with the lottery and the preface paradox. Previously it was thought that *Solutions from Statistical Evidence* or *Solutions from Defending the Conjunction* would be all that is needed to defend closure against the pressure coming from closure puzzles. However, it now seems that that these solutions have targeted features of the respective puzzles that turn out to be non-essential for constructing closure-threatening paradoxes. In other words, the features that these popular closure-preserving solutions target, rather than being necessary features of closure-threatening paradoxes, are merely incidental features of the lottery and the preface narratives. What the Paradox of the Pill illustrates, is that in close vicinity to the lottery and the preface paradox there lurk

relevantly similar epistemic puzzles to which the prominent closure preserving solutions that were previously thought to get closure out of trouble do not apply.

The central idea underlying *Solutions from Statistical Evidence* put forth in response to the lottery paradox was that purely statistical evidence, no matter how strong, does not justify beliefs. This strategy identifies the statistical nature of the evidence as a crucial feature of the lottery paradox. Thus, proponents of *Solutions from Statistical Evidence* identify something along the lines of Nelkin's p-inferences as an essential feature of the puzzle.

(P-inference) p has a statistical probability of n [where n is a very high number] \rightarrow p.

Subsequently, proponents of closure proposed modifications to the Sufficiency Thesis that would exclude beliefs based solely on statistical evidence from being justified. Recall Nelkin's proposal that for a belief to be justified an agent must be in a position to *presuppose* a causal connection to the facts. This strategy would of course be successful if it was an essential part of the puzzle that the individual conjuncts be supported by purely statistical evidence. What Nelkin and her followers seem to have overlooked however, is that one can be in a position to presuppose a causal connection to the evidence but still be in a lottery like environment. In the pill scenario this is achieved by building a lottery mechanism into S's causal connection to the evidence. As a result, we can generate challenges to the closure principle without relying on evidence that is purely statistical. In other words, we can generate challenges for (MPC) without relying on p-inferences. This suggests that *Solutions from Statistical Evidence* are not sufficiently general to achieve their aim of saving (MPC) as they target a non-essential feature of the lottery-paradox. The statistical nature of the evidence, rather than being an essential feature of closure-threatening paradoxes is merely an incidental feature of the lottery narrative; in this sense the statistical nature of the evidence is a red herring.

Furthermore, strategies like *Solutions from Statistical Evidence* that set out to solve the lottery paradox by rejecting (L1) - the premise that we can justifiably believe that any given lottery ticket will lose - benefit from the fact that denying justification in lottery propositions is intuitively acceptable. In other words, denying justification

for lottery beliefs does not strike us as particularly problematic or implausible. A possible explanation for this might be that when it comes to lottery tickets, we generally aren't overly concerned about the exact epistemic status of our beliefs. If it turned out that we are not justified in believing lottery propositions, then so be it; especially if this allows us to preserve the attractive principle of (MPC). After all it is plausible that our intuitive pull towards (MPC) outweighs our concerns over the exact epistemic status of lottery beliefs. However, things are different in the case of The Paradox of the Pill. In this case rejecting (BP1), i.e. denying that any of S's beliefs are justified, no longer strikes us as a plausible or acceptable solution. While it might have been intuitively acceptable to deny justification for individual lottery beliefs, it does not seem plausible to respond to a miniscule amount of risk (0.01%) by embracing global skepticism about justification. It seems again that solution to the lottery paradox that reject (L1) rely on a specific feature of the paradox - in this case the intuitive acceptability of denying justification in lottery beliefs – that is not an essential or general feature of closure puzzles but instead only a by-product of the particular lottery narrative. The Paradox of the Pill shows that we can create cases in which denying justification for the individual conjuncts is no longer a plausible strategy for preserving (MPC).

A similar worry arises for *Solutions from Defending the Conjunction* produced in response to the Paradox of the Preface. These solutions, as proposed by Kaplan for instance, pivot on the idea that in preface cases the author lacks justification for believing that the conjunction in question is false as there is no 'special reason' supporting it. Subsequently it is suggested that we replace (P3) with the weaker premise (P3*).

(P3*) The author is justified in believing that it is *likely* that errors will be found in her book.

In order for this strategy to apply more generally however, Kaplan and his followers must assume that the lack of a special reason against the conjunction is an essential feature of closure-threatening puzzles. However, as The Paradox of the Pill shows, we can easily construct closure-threatening paradoxes that avoid this feature of the

preface paradox. In the paradox of the pill there does exist a special reason to believe that a sufficiently long conjunction is false, it is guaranteed by the pill's effects. Subsequently we cannot replace (BP3) with a weaker premise (BP3*).

(BP3*) S is justified in believing that the conjunction of her justified beliefs will *likely* contain at least some false beliefs

(BP3*) would fail to correctly capture the epistemic environment the pill creates, as the pill does more than just make it likely that a sufficiently long conjunction will be false – it guarantees it. Thus, *Solutions from Defending the Conjunction* also seem to target a feature of the preface paradox that is merely incidental to the preface narrative but not essential for the creation of closure-threatening puzzles. As a result, this strategy also fails to provide a more general solution to the paradoxes that threaten (MPC).

What the Paradox of the Pill shows, is that we can create problematic puzzles that do not incorporate any of the features previously targeted by many of the popular closure-preserving solutions to the lottery paradox and the paradox of the preface. This is the case because both strategies have targeted features of the respective puzzle that are non-essential for creating puzzles that put pressure on (MPC). In light of this generality worry, it seems that Lockeans who want to preserve closure will need to revise their current strategies or explore alternatives if they want to solve the pill paradox whilst retaining multi premise closure. Until then however the rejection of (MPC) seems to remain the most promising solution to the new paradox.

8. *Problems From a Different Direction*

At this stage Lockeans who are sympathetic to (MPC) might offer the following concessionary response: Perhaps we were too quick to make high probability both *necessary* and *sufficient* conditions for justification. In other words, perhaps we were too quick to formulate the Lockean View as a biconditional. What the Paradox of the Pill does shows is that there are a good reason to revise our view - perhaps high probability is necessary but *not* sufficient for justification. So, to preserve (MPC) we

should revise the formulation of the Lockean View by restricting it to the following Necessity Thesis and remain silent on justification's sufficiency conditions.

Necessity Thesis (NT) The belief that P is justified for S, only if P is highly probable for S.

If the Lockean View is restricted in this way, then, one might hope, the pressure on closure, which the Paradox of the Pill exerts can be avoided. A reason for thinking this is that the paradox explicitly relies on the Sufficiency Thesis - i.e. the idea that high probability is sufficient for justification. However, this argument turns out to be unsuccessful.

First, we might worry that the Lockean retreat of denying the sufficiency direction of the Lockean View amounts to a significant departure from their probabilistic project. There seems to be a tension between thinking that justification is best explained in terms of high probability or sufficient likelihood while at the same time conceding that some beliefs that are highly probable are not justified. By denying the sufficiency direction of the Lockean View, Lockeans effectively concede that there are problems with the probabilistic conception of justification and this threatens to undermine the overall appeal of their view. To make this point more succinct, we might worry that at this point Lockeans are conceding too much. However, for now we must not worry about these admittedly very general concerns, for there is a much more immediate reason to think that the retreat to the Necessity Thesis is not going to help with the aim of preserving (MPC).

Even if we restrict the Lockean View to the Necessity Thesis, (MPC) still fails due to the problem of *risk agglomeration*. Consider a number of uncontroversially justified beliefs (e.g. beliefs based on perception, testimony, or memory) that are individually highly probable but not *necessarily*⁹³ true for S. In other words, consider a number of individually justified beliefs such that P's degree of probability for S is above the threshold required for justification but below 1. Now, as we begin

⁹³ By this I simply mean that it is possible for the beliefs to be false or that they have a non-zero possibility of being false. This will be the case for any belief with the exception of beliefs in logical truths or mathematical truths, which are standardly taken to be true in all possible worlds.

conjoining these individually justified beliefs the error-risk associated with each conjunct accumulates until eventually the probability of the conjunction falls below the threshold required for justification. So, we can end up with cases in which conjoining a number of highly probable beliefs yields a highly improbable, and therefore unjustified, conjunction. As a result, (MPC) will fail on the Lockean View even if we restrict it to the necessity direction of the original biconditional. So, in the end, there is little hope that Lockeans will be able to avoid the objection that their view is incompatible with (MPC).

9. *A Problem that Sticks*

While the previous chapter focused on arguments against the Lockean View that turned out to be unpersuasive, this chapter has focused on a challenge that is ultimately more successful. It has been shown that the Lockean View is incompatible with the compelling and intuitively attractive principle of multi premise closure - a principle that we have argued is a plausible desideratum for a satisfactory theory of justification.

I argued that problems for the viability of (MPC) come from two directions – the sufficiency direction as well as the necessity direction of the Lockean View. From the sufficiency direction of the Lockean View – we called it the *Sufficiency Thesis* (ST) - combined with (MPC) and (NC) we can derive a forceful new paradox, *The Paradox of the Pill*. As we have seen, the most promising solution to this paradox seems to be the rejection of (MPC). At the same time there is pressure on (MPC) from the necessity direction. Even if we restrict the Lockean View to the Necessity Thesis, (MPC) will still fail due to the problem of risk agglomeration.

Thus, with pressure from both direction of the biconditional, it appears that there is little hope that Lockeans will be able to avoid the objection that their view is incompatible with (MPC). Unlike, the challenges from statistical evidence considered in the previous chapter, the multi premise closure challenge then seems to be an objection that sticks. And since we identified (MPC) as a plausible desideratum for a satisfactory theory of justification, we have a reason to be unhappy with the Lockean View and by extension any account of justification that follows it.

At this stage the following thought seems to arise almost naturally: If probabilistic approaches to justification like the Lockean View are incompatible with (MPC), then perhaps relocating the source of justification to something other than high probability will allow us to avoid this problem. Putting things in terms of properties, we may express this thought as follows. If identifying the property of *being highly probable* as the justification-conferring property forces us to deny (MPC), then perhaps finding a different, *non-probabilistic*, justification-conferring property will put us in a position to preserve (MPC). In the next chapter I consider the prospects of such non-probabilistic alternative.

Chapter 4

A Non-Probabilistic Alternative: Normalcy Views and the Problem of Easy-Defeat⁹⁴

1. Towards a Non-probabilistic Approach to Justification

Up to this point we have introduced the Lockean View of justification – the perhaps most widely endorsed picture of epistemic justification – and considered a number of objections. According to the Lockean View, a belief that P is justified for S iff P is highly probable for S.

Lockean View A belief that P is justified for S iff P is highly probable for S.

The perhaps most serious objection to the Lockean View, as was argued in the previous chapter, is that it is incompatible with the intuitively compelling and attractive principle of multi premise closure.

Multi Premise Closure (MPC) If S is justified in believing p and S is justified in believing q... and S is justified in believing n, then S is justified in believing the conjunction (p & q... n).

The case against the compatibility of Lockean Views with (MPC) comes from two directions. The sufficiency direction of the Lockean View combined with (MPC) and a plausible No Contradiction principle gets us into a number of paradoxes, which are most plausibly solved by rejecting (MPC), while the necessity direction gives rise to

⁹⁴ This chapter is based in large part on my article “Normalcy, Justification, and the Easy-Defeat Problem” forthcoming in *Philosophical Studies*.

the problem of risk accumulation; as one begins conjoining individually justified beliefs, the error risk associated with each conjunct accumulates and eventually the probability of the conjunction can fall below the threshold required for justification. These challenges provide good reasons for thinking that on probabilistic accounts of justification (MPC) will fail. However, giving up (MPC), as we have seen, is not just counterintuitive but also comes at a serious theoretical cost - if we deny (MPC), then we can no longer use conjunction introduction to generate new justified beliefs from old ones.

Motivated primarily by the incompatibility of probabilistic accounts with (MPC), epistemologists have become increasingly interested in exploring alternative accounts of justification. The guiding thought is the following: if probabilistic accounts are not compatible with closure, then maybe relocating the source of epistemic justification to something other than high probability will allow us to avoid this problem.

One strategy for constructing non-probabilistic accounts of justification would be to simply strengthen the Lockean View. Instead of requiring that a belief be *highly probable*, we could instead insist that in order to be justified a belief must have probability 1. What exactly this amounts to would of course again depend on what one takes the *epistemically relevant good-making feature* to be. For internalists, for instance, this would mean that one's level of evidential support would have to be 1 – i.e. in order to be justified the evidence would need to entail a proposition. Alternatively, for simple *process reliabilist*, it would mean that in order for a belief to be justified it would have to be formed using a belief-forming method that is perfectly reliable across the entirety of modal space. For our current purposes however, the general formulation in terms of probability 1 will suffice. In modal terms, this strengthened version of the Lockean View can be expressed in what we may call the *Strong View*.

Strong View A belief that P is justified for S iff P is true in all possible worlds.

While the Strong View would be successful at avoiding the problem of risk accumulation – after all in order for any belief to be justified it must be true in all

possible worlds – it faces an obvious objection. By requiring that a proposition be true across the entirety of modal space the Strong View becomes much too demanding.

According to the received view, only propositions expressing *necessary truths* (e.g. logical and mathematical truths) are true in all possible worlds. An important feature of necessary truths – since they are true in all possible worlds – is that they couldn't possibly be false. However, most of our beliefs are not about logic or mathematics. Rather, many of the beliefs we ordinarily care about are beliefs in propositions that, if true, are only *contingently* true – i.e. they are beliefs in propositions that have a non-zero chance of being false. Put differently, most of the beliefs that we ordinarily care about are such that, unlike beliefs in necessary truths, they *could possibly be false*. For instance, for beliefs based on perception, memory, or testimony, there always exists a possible world in which these beliefs are false; such worlds include worlds in which we are deceived by an evil demon, worlds in which our memory malfunctions, or worlds in which a testifier is trying to mislead us. As a result, on the Strong View, no contingent proposition could be justifiably believed and subsequently – since most propositions we care about in day-to-day life are contingent in nature - we would be justified in believing hardly anything at all.

Moreover, since in order to be justified a belief cannot possibly be false, accounts of justification following the Strong View will fail to satisfy at least one of justification's core platitudes, namely that justification is fallible.⁹⁵ And since the platitudes provide constraints that any theory needs to satisfy in order to be a theory about epistemic justification rather than something else, we should be reluctant to accept accounts following the Strong View as accounts of justification. After all it is incompatible with at least some of the principles that are widely believed to be essential to the notion of epistemic justification. In short, according to the Platitudes-First Approach outlined in chapter 1, the Strong View would fail to be a theory of justification.⁹⁶

⁹⁵ See Chapter 1, Sect. 3.1.

⁹⁶ It may of course still be an interesting theory of something else. But it would not count as a theory of justification.

Fortunately, more promising, non-probabilistic, accounts of justification have recently been proposed by Martin Smith (2010, 2016) and Jarrett Leplin (2009). Instead of considering a belief's modal profile across *all possible worlds*, these accounts are less demanding as they constrain the epistemically relevant worlds to the set of *normal worlds*. On these views, broadly speaking, in order for a belief that P to be justified for S, it is necessary that given the evidence, there exists no *normal world* in which S falsely believes that P. Let's call this picture of justification the *Normalcy View*.

Normalcy View (NV) The belief that P is justified for S only if, given S's evidence E, there does not exist a single *normal world* in which S falsely believes that P.

According to the Normalcy View the justification-conferring property - i.e. the J-property - is some property along the lines of *being true in all normal worlds*; or, for short of *being normally true*. As we will see, depending on how one spells out the details of (NV) one will end up with slightly different J-properties, but ultimately, on Normalcy Views, the relevant J-property will amount to something in the vicinity of *being normally true*.

Like the Strong Thesis, the Normalcy View is able to preserve (MPC) as it avoids the problem of risk accumulation: since in order for a belief to be justified there can exist no normal world in which the belief is false, there will also be no normal world in which a conjunction of these beliefs is false. After all, a normal world in which the conjunction is false would require a normal world in which one of its conjuncts is false and this possibility is ruled out by (NV). However, unlike the Strong View, the Normalcy View does not require that a belief must be true in all possible worlds in order to be justified or that it must be impossible for the belief to be false. As a result, unlike the Strong View, the Normalcy View is compatible with fallibilism - the idea that there can be justified false beliefs.

The Normalcy View is compatible with justified false beliefs as long as the worlds in which S falsely believes that P are not amongst the set of normal worlds. In other words, (NV) is compatible with justified false beliefs in cases where the falsity of the belief is explained by abnormal circumstances. This seems to put Normalcy

Views in the unique position of being able to accommodate (MPC) without requiring a commitment to infallibilism. So, the Normalcy View is considerably less demanding than the Strong View.

Hence, the Normalcy View seems to be a *prima facie* promising alternative to the probabilistic approach to justification underlying the Lockean View and it is the general prospects of this new family of views that is the primary focus of this chapter.

2. *The Normalcy View and Justification's Core Platitudes*

Before considering the two leading instances of the Normalcy View – Smith's (2010, 2016) *Normic Support Account* and Leplin's (2009) *Normic Reliabilism* – it will be helpful to first consider whether Normalcy Views will be able to accommodate justification's core platitudes. Recall that in order to be a theory of justification rather than something else, a theory needs to satisfy justification's core platitudes. Also recall that justification's central platitudes proposed in chapter 1 were *Truth Candidacy*, *Permission*, and *Blamelessness*. One might wonder whether answering any questions about the Normalcy View requires that we first clarify what it means for a world to be *normal*. However, while this will play an important role later, for the task at hand these details do not matter, and we can rely on an intuitive notion of normality (if it helps, we can, for now, simply treat *normal worlds* analogous with *close worlds*).

Truth Candidacy. According to the Truth Candidacy platitude, a theory of justification needs to somehow accommodate the thought that a belief is justified *only if* it is a good candidate for being a true belief. In other words, any candidate justification-conferring property must be such that it does not apply to beliefs unless they are good candidates for being true. Can we offer a plausible explanation for why the Normalcy View meets this constraint? I think we can. While Lockeans spelled out the notion of being a *good candidate* for being a true belief in terms of high probability, proponents of Normalcy Views may argue that *being true in all normal worlds* is also a property that makes a belief a good candidate for being a true belief. After all, it seems convincing that a belief which is true in *all* normal worlds and would only be false under abnormal circumstances should count as a good candidate

for being true. Hence, the Normalcy View can plausibly accommodate the Truth Candidacy platitude.

Permission. According to the Permission platitude, if S has justification for believing P, then it must be *at least* permissible for S to believe that P. Can we make sense of this platitude on the Normalcy View? Again, I think we can. On the Normalcy View, in order to be justified, a belief must be true in all normal worlds given one's evidence. Put differently, on normalcy views justified beliefs will only be false if abnormal conditions obtain. Proponents of the Normalcy View can use this idea to explain why Permission holds on their view: if one's evince is such that one would only falsely believe that P under abnormal conditions, then surely it should be permissible for S to believe that P.

Blamelessness. Finally, we need to consider whether the Normalcy View can accommodate the platitude that if a justified belief turns out to be false, then the agent will be free from blame. Again, it seems that they plausibly can. The story will be similar to the previous one about Permission. If a justified belief turns out to be false on the Normalcy View, then, since justification requires that a belief be true in all normal worlds, it must be the case that abnormal conditions have obtained. However, since we are not in control of our external environment, it seems reasonable that a subject should be blameless in cases where the falsity of their belief is due to abnormal conditions. There are limits to what we can reasonably expect from a subject and if it is the case that given their evidence, the belief that P is true in all normal worlds, it seems reasonable to think that they have done what can reasonably be expected of them. So, the Normalcy View seems to have no problem satisfying the Blamelessness platitude.

The Normalcy View then does not just offer the promise of preserving (MPC), but unlike the Strong View, it also satisfies justification's core platitudes. This is a very promising result.

The remainder of this chapter aims to raise some trouble for normalcy views by arguing that they give rise to problematic notions of epistemic defeat. In the next section I will briefly present the two leading instances of the Normalcy View: Smith's *Normic Support Account* and Leplin's *Normic Reliabilism*. I then argue that both accounts face the same problem. Both accounts are highly sensitive to defeating

evidence, which results in justification being lost much too easily - I will call this the *Easy-Defeat Problem*. Finally, I argue that since the problem is structural in nature there is little hope that other views following (NV) will be able to avoid this problem. In the end it appears that there are good reasons for rejecting the central claim of the Normalcy View, namely that in order for a belief to be justified it is *necessary* that there does not exist a single normal world in which S falsely believes that P.

3. *The Two Leading Proposals*

As presented above (NV) only provides a very general modal constraint on theories of justification. As such it only provides the general contours, or the modal shape, of more complete theories. The details of how to best understand the modal constraint proposed by (NV) can of course be developed in different ways. There are currently two leading proposals on the market.

The first one is due to Martin Smith who recently proposed an account of justification according to which a belief that P is justified only if the evidence E *normically supports* P - i.e. if given one's evidence, P is true in *all* normal worlds. Let's call this the *Normic Support Account*.

Normic Support Account In order for one to have justification for believing a proposition P, it is necessary that one's body of evidence E normically support P – it is necessary that all the most normal worlds in which E is true are worlds in which P is true. (2016, p. 42)⁹⁷

As we can see, in the Normic Support Account the modal constrained presented by (NV) is cashed out via a specific notion of evidential support called *normic support*.

⁹⁷ Formally, Smith expresses the account as follows. Let $E \blacksquare \rightarrow P$ represent 'E normically supports P' and let \mathcal{N}_w be a function "carrying information about the comparative normalcy of possible worlds from the perspective of w" (Smith 2016: 137). Then the notion of normic support can be expressed as follows:

Normic Support Account Formal. $E \blacksquare \rightarrow P$ is true at $w \in W$ iff either (i) there is a sphere $N \in \mathcal{N}_w$ such that N is E-permitting and every E-world in N is a P-world or (ii) there is no sphere $N \in \mathcal{N}_w$ such that N is E-permitting. (137)

A very different way of developing (NV) can be found in recent work by Jarrett Leplin. According to Leplin a belief is justified only if it was *reliably produced*; where reliably produced means produced using a belief-forming method that is *perfectly* reliable under normal conditions. More precisely, for Leplin a belief is justified only if (i) it was produced by a method that is *perfectly* reliable under normal conditions and (ii) the believer has no reason to believe that conditions are abnormal (2009: 43).⁹⁸ In spelling out what it means for a belief-forming method to be perfectly reliable, Leplin draws on modal considerations inspired by Nozick’s notion of sensitivity. According to Leplin, “[a] method of belief-formation is reliable if it would not produce or sustain false beliefs under normal conditions. If a belief is produced or sustained by a reliable method under normal conditions, then were this belief to have been false, the process would not, under those conditions, have produced or sustained it” (35). For our purposes, Leplin’s account of justification (henceforth *Normic Reliabilism*) might be presented as follows.

Normic Reliabilism In order for S to have justification for believing a proposition P, it is necessary that S’s belief was reliably produced – it is necessary that (i) the belief that P was produced by a belief-forming method M that would not, in any normal world, produce the belief that P if P were false and (ii) that the believer has no reason to believe that conditions are abnormal.

We can see that there are important differences in how the two leading Normalcy Views develop (NV). While the Normic Support Account accommodates the modal constraint proposed by (NV) via a special notion of *evidential support*, Normic Reliabilism accommodates (NV)’s modal constraint via a particular notion of *reliable belief-production*.⁹⁹ But, we will see that despite these differences both accounts turn out to be problematic.

⁹⁸ Leplin argues for a third condition, namely that the belief-forming method must be used *intentionally*. This condition is intended to address the well-known Generality Problem for reliabilist accounts of justification. Since I am not concerned with this issue here I will, in the interest of simplicity, ignore this detail.

⁹⁹ Interestingly, despite the considerable differences between Smith’s Normic Support Account and Leplin’s Normic Reliabilism it is very difficult to think of a case in which the two theories make

In the next two sections I argue that the Normic Support Account is unsatisfactory because it gives rise to an account of epistemic defeat on which beliefs lose their justificatory status much too easily. I then argue that the same problem befalls Normic Reliabilism.

4. Normic Support

According to the Normic Support Account a belief is justified only if it is *normically supported* by the evidence. What does it mean for a body of evidence to normically support a proposition?

Informally, a body of evidence E normically supports a proposition P only if, given the evidence E, it would call for *special explanation* if P turned out to be false (Smith 2016: 40). To illustrate, consider some paradigmatic instances of justified beliefs. For example, suppose a subject wrongly forms a belief based on perceptual evidence (e.g. that the car in front of them is red). Smith suggests that this case would call out for a *special explanation*; some mitigating or interfering circumstances must have obtained to explain why a belief based on perception would fail to be true. Such circumstances might include hallucinations, unfortunate lighting, or deception by a malevolent demon, etc. Similar considerations will of course apply to beliefs based on testimony or memory. Cases in which a belief based on testimony turns out to be false would require some sort of special explanation, e.g. intentional deception by the testifier, a testifier misremembering the facts, mishearing by the receiver, etc. Informally then, for a belief to be normically supported by the evidence, it has to be the case that if the belief turned out to be false, some special explanation would be required. Importantly, the ‘calling for an explanation’ relation is supposed to be a *metaphysical* relation between a body of evidence E and a proposition P. In other words, it is supposed to be an objective –i.e. mind independent – relation between a

diverging predictions. As an interesting working hypothesis, it might be suggested that the two theories are extensionally equivalent.

body of evidence and a proposition. It is *not* supposed to be a *psychological* notion that simply tracks what we find surprising or curious.¹⁰⁰

To turn the notion of normic support into a more formal notion, Smith ties the notion of ‘calling for an explanation’ to the notion of *normality*. Events that are normal, so the thought goes, do not call for special explanations, while events that are abnormal do call for special explanations. This allows Smith to provide an analysis of ‘normic support’ in terms of *normal possible worlds*.¹⁰¹ Beliefs that are normically supported by the evidence will *normally* be true - i.e. they will be true in normal worlds. Smith captures this idea in the following modal account of normic support:

Normic Support Modal. A body of evidence E normically supports a proposition P just in case P is true in all the most normal worlds in which E is true (Smith 2016: 42).

102,103

This account of normic support combined with the claim that a belief is justified only if it is normically supported by the evidence yields the full Normic Support Account,

¹⁰⁰ This allows the account to avoid worries surrounding the idea that there will be variation across different individuals regarding what does and what does not require a special explanation. For instance, someone who does not care about a certain proposition or has a very low level of curiosity may be less inclined to ask for special explanation compared to someone who is generally very curious about things or someone who is disposed to be very easily surprised. Since the ‘calling for an explanation’, on Smith’s account, is a *metaphysical* relation between a body of evidence and a proposition these *Inter-Personal Variation Worries* will have no purchase on the Normic Support Account, as they do not affect the notion of normic support and subsequently do not affect whether a proposition is justified or not.

¹⁰¹ More precisely, ‘calling for an explanation’ becomes a function carrying information about the comparative normalcy of possible worlds.

¹⁰² This modal account of normic support is of course reminiscent of certain well-known safety conditions on knowledge. A relevant asymmetry however is that in the normalcy framework worlds are not ordered in terms of *comparative similarity* but instead in terms of *comparative normalcy*. One important difference between ranking worlds in terms of their comparative normalcy as opposed to their comparative similarity is that the former *does not* require centering on the actual world (strong centering) while the latter does. In other words, while no world will ever be more similar to the actual world than the actual world itself, the actual world must not always be the most normal world; for sometimes abnormal things happen in the actual world.

¹⁰³ In order to get the modal account of normic supports off the ground, we need to grant that worlds can be ranked according to their *comparative normalcy*. Smith (2016: 42) is explicit about this assumption and I see no reason to challenge it here.

according to which a belief in P is justified for S only if, given S's evidence E, there exists *no* normal world in which E and not-P obtain together. More informally we might say that one is justified in believing P only if one's evidence does not contain any information that would explain why one might falsely believe that P.

4.1. Two Initial Worries

At this point there are two potential worries for the Normic Support Account that we should address.

The first worry is that the notion of 'calling for an explanation' is too imprecise to do the work it is supposed to do. Consider that in order for the Normic Support Account to be successful we need to be able to compare worlds in terms of their respective degrees of normality. This is supposed to be done via the notion of 'calling for an explanation'. But what determines whether something calls for an explanation or not? Recall that the 'calling for an explanation' relation is supposed to be a *metaphysical* relation between a body of evidence and the world, and not a *psychological* one. But how could we determine whether something *metaphysically* calls for an explanation? In a recent review of Smith's (2016) book, Anderson (2017) put pressure on this issue, "The main shortcoming of the book is that the core notion is not given sufficient analysis. Normic support is, *prima facie*, an intuitive notion. Unfortunately, Smith leaves too much to our intuitive grasp...a detailed account of what the notion could do for us is inadequate without a solid grasp of the notion itself." Anderson's worry can be sharpened further by pointing out that if indeed too much is left to our intuitions, then the Normic Support Account runs the risk of collapsing into a psychological account, which Smith is trying to avoid (see footnote 100). In short, one might worry that Smith does not tell us enough about how to determine whether a world is metaphysically normal or not.

While I take this to be a good objection, here is how proponents of the Normic Support Account may respond. It is widely accepted that for any given world we can determine what the *close* or *relevantly similar* possible worlds are - accounts of sensitivity (Nozick 1981), safety (Sosa 1999; Williamson 2000), and epistemic luck (Pritchard 2005) all rely on this assumption. However, it is not like we have a more

precise or more principled method for ranking worlds according to their comparative similarity or closeness. So, the problem of not having a principled method for comparing worlds for their comparative normalcy does not seem to be a problem specific to the Normic Support Account. Instead it seems to be a very general problem befalling many epistemic theories.

The second worry is slightly different in nature - it concerns predictions (or the extension) of the Normic Support Account rather than particular features of how the view is formulated. One important feature of the Normic Support Account, since it does not explain justification in probabilistic terms, is that whether a belief that P is normically supported by the evidence is logically independent from P's degree of probability. In other words, the two notions – *normic support* and *high probability* – come apart. This turns out to be less mysterious than it may initially appear. Consider the belief that one's ticket in a very large lottery is going to lose. While it is highly probable that this belief is true, it wouldn't call for a special explanation if it turned out to be false – i.e. if one happened to win the lottery. As a result, beliefs about lottery tickets, and other beliefs based on purely statistical evidence, will always fail to be justified on the Normic Support Account. However, as we pointed out in chapter 2, many do judge that in cases of large lotteries we are justified in believing that one is not going to win. We may consider this a strike against the Normic Support Account. But, ultimately, I do not think that this is a particularly forceful objection against the account – as I argued in the previous chapter, while it is perhaps counterintuitive to deny that lottery beliefs can be justified, it will nevertheless strike many as ultimately *acceptable*; especially if this will allow us to preserve (MPC). So, I do not think that proponents of the Normic Support Account will be particularly impressed by this concern over the justificatory status of lottery beliefs. Hence, I will put this issue aside for now and not pursue it further here. It is important to note however these considerations point towards a more general consequence of the Normic Support Account, namely that it is possible for a belief to be overwhelmingly probable on the evidence and yet fail to be justified.¹⁰⁴ And this more general issue I do want to pursue

¹⁰⁴ Likewise, it is possible for a belief to be highly improbable on the evidence and yet to be normically supported by the evidence. One example comes from Smith's analysis of the preface paradox. Smith argues that the author's belief that the book is error-free, despite being highly improbable, can

further. As we will see this insensitivity to probabilistic considerations will turn out to be problematic for the Normic Support Account.

5. *The Easy-Defeat Problem*

5.1. An Initial Skeptical Worry

One might worry that denying the justificatory status of lottery beliefs, or beliefs based on statistical evidence more generally, exposes the Normic Support Account to a skeptical threat. After all, haven't Vogel (1990, 1999) and Hawthorne (2004) given us compelling reasons to think that many of the propositions we believe entail lottery propositions? If lottery propositions fail to be justified on the Normic Support Account, and Hawthorne and Vogel are correct that many ordinary beliefs entail lottery propositions, then the Normic Support Account is threatened to be undermined by skeptical worries.

To make this worry more concrete, consider the following examples.

- (a) That Donald Trump is the current president of the United States entails that Donald Trump did not suffer from a fatal heart attack within the last few minutes.
- (b) That I will be able to cook dinner later entails that there will not be a power outage in my neighborhood.
- (c) That I am correct in believing that the New York Knicks beat the Boston Celtics last night after reading it in the newspaper entails that the newspaper did not misprint the scores.

Note that these entailments appear relevantly similar to lottery propositions in that they are (a) statistically very likely but (b) don't have any special evidence in their support. Since the Normic Support Account denies that high probability is sufficient for justification, one might worry that it will end up having to deny that we can

nevertheless be normically supported by the evidence. We will look at this issue more closely in chapter 6.

justifiably believe that Donald Trump is the current president of the United States, that one will be able to cook dinner later, or that the New York Knicks beat the Boston Celtics last night. Similar lottery-like entailments can of course be found for almost any propositions. How can proponents of the Normic Support Account respond to this skeptical challenge?

Fortunately, Smith offers what I take to be a compelling response. In short, he tries to break the supposed symmetry between these *lottery-like* entailment proposition and *genuine* lottery propositions. The thought is the following: while for genuine lottery beliefs, e.g. the belief that my lottery ticket is going to lose, it would not call for a special explanation if the belief turned out to be false, for *lottery-like* propositions it would call for a special explanation if they were false.

“My response to the skeptical problem is perhaps not as definitive as one might wish – but I want to suggest that the analogy between genuine lottery propositions and lottery-like propositions is simply not as close as Hawthorne suggests.... If the president died within the last five minutes, or there is a power outage in my neighborhood, or the scores in the newspaper are in error, there would have to be some explanation as to how such things came about... This is a significant disanalogy with genuine lottery propositions” (Smith 2016: 57-58).

A different way of making Smith’s point is to say that while for *genuine* lottery propositions, given the evidence, there always exists a normal world in which one’s belief is false, for *lottery-like* propositions such as the ones considered above, there does not exist a normal world in which they are false. As a result, one would lack justification for believing that one’s lottery ticket is a loser but one would be justified in believing that Trump has not recently suffered a fatal heart attack and that he is the current president of the United States. Thus, the skeptical threat of traditional Hawthorne/Vogel lottery variants is avoided.

However, in what follows I argue that the notion of defeat that falls out of the Normic Support Account gives rise to a new skeptical challenge – one which is not so easily avoided.

5.2. The Easy-Defeat Problem.

Recall that according to the Normic Support Account, a belief that P is justified only if the evidence E normally supports P, i.e. only if *all* of the normal worlds in which E is true are worlds in which P is true. It follows that justification is defeated as soon as a piece of evidence establishes the existence of *a single* normal E & not-P world. In these cases, it would no longer call for an explanation if the belief turned out to be false, because a world in which P is false is compatible with one's body of evidence. From this we can construct the following account of *Normic Defeat*.

Normic Defeat (ND) S's justification for believing P is defeated if given a new piece of evidence D, there exists at least one normal world in which E is true and P is false.

(ND) is a non-probabilistic account of defeat and the fact that non-probabilistic accounts of justification will generate non-probabilistic accounts of defeat is not surprising. However, non-probabilistic accounts of defeat like (ND) do have surprising consequences: since the Normic Support Account is entirely insensitive to probabilistic considerations there is nothing that prevents *minimal probability possibilities*, i.e. possibilities that are overwhelmingly improbable, from acting as defeaters. As soon as a piece of evidence establishes the existence of just *one* normal E & not-P world, no matter how remote the probability that this world is the actual world, justification for P is lost.

In this regard, the normic support framework departs considerably from that of standard threshold accounts on which defeating evidence needs to be sufficient to bring a belief's degree of probability below the threshold required for justification. According to (ND) however, for a piece of evidence to defeat P's justificatory status it is not necessary that it makes the possibility of being in an E & not-P world even remotely probable; all it needs to do is introduce *a single* E & not-P world into the sphere of most normal E worlds.

This should provide some initial reason for concern. If all that is required for justification to be defeated is for some evidence to introduce the existence of just *one*

normal E & not-P world, no matter remote the probability that this world is the actual world, then one might worry that defeat simply comes too easily on the normalcy view. This problem - let's call it the *Easy-Defeat Problem* - has some unattractive and costly consequences.

We can illustrate the Easy-Defeat Problem by slightly *augmenting* traditional Vogel/Hawthorne style lottery variants.

Allergy Helen is allergic to peanuts. She goes to a café and orders a brownie labelled 'peanut free'. Based on this Helen is justified in believing P, that the brownie is safe to eat. On the Normalcy View this entails that there does not exist a single normal world in which Q, the brownie has been contaminated by something containing peanuts. After ordering the brownie Helen sees a newspaper headline that reads D, international flour supplier admits to having accidentally put one bag of peanut-contaminated baking flour into circulation.

Allergy differs from standard lottery-like cases in that it does not just raise to salience the plain possibility of error, i.e. that there is always a small chance that even a brownie labelled 'peanut-free' might turn out to be contaminated. Instead it goes one step further and provides, through additional evidence D, an explanation for why Helen's belief that her brownie is safe to eat may be false. It is this additional piece of evidence that makes it an *augmented* case of its more traditional counterparts. I assume that in Allergy we would still want to say that Helen is justified in believing that her brownie is safe to eat, and that D is not sufficient to defeat the justificatory status of Helen's belief; after all we do not ordinarily assign such seemingly negligible evidence this kind of defeating power. However, as I will argue now, the Normic Support Account will not deliver this result.

In order for Helen to be justified in believing P, that her brownie is safe to eat, there must not exist a single normal world in which Helen, given her total body of evidence, falsely believes that P. Since a world in which Q is true would be such a world, in order for Helen's belief in P to be justified there cannot exist a single normal world in which Q obtains, i.e. there cannot exist a single normal world in which Helen's brownie has been contaminated. An important question then is whether given

Helen's total body of evidence E, which now includes D, the information about the one contaminated bag of flour, there exists at least one normal world in which Helen's brownie is contaminated.

It seems plausible that such a world does exist – this of course is the world in which the one bag of contaminated flour just so happened to end up in the café at which Helen just ordered her brownie. Given Helen's evidence there seems to be nothing strange or abnormal about this possibility. In the end the bag of contaminated flour must have ended up somewhere and there does not appear to be anything particularly abnormal about it having ended up in the café where Helen just bought her brownie than in any other café. It might of course be *less probable* that it would end up in some café rather than others – after all, some cafes will go through a lot more flour than others – but recall that the normalcy account ignores probabilistic considerations of this kind. So, given Helen's total body of evidence it seems plausible that there *does* exist a normal world in which Helen's brownie has been contaminated. As a result, Helen's belief that her brownie is safe to eat fails to be normically supported by the evidence and subsequently fails to be justified on the Normic Support Account. A problematic prediction.

It is of course easy to generate more of these cases. Consider the following example.

Lightning A few days ago Helen has made plans with her friend Bob to visit him in Oxfordshire next weekend. Based on this Helen is justified in believing P, that she will see Bob next weekend. On the Normalcy View this entails that there does not exist a single normal world in which Q, Bob has been fatally struck by lightning. As Helen is thinking about her upcoming trip to Oxfordshire she reads a newspaper headline stating D, man in Oxfordshire fatally struck by lightning.

The worry is the same as in the previous case. In light of D, a seemingly negligible piece of evidence, it appears that Helen's belief is no longer normically supported by the evidence. Why? Consider the following question: Given Helen's *total evidence*, does there exist at least one normal world in which her belief that she will see Bob next weekend is false? Again, it seems plausible that such a world exists - this of

course is the world in which the man fatally struck by lightning just happened to be Bob. Given that Helen *knows* that a man was fatally struck by lightning, there does not appear to be anything terribly abnormal about the possible world in which this person was Bob. Again, it is of course highly unlikely that Bob was struck by lightning but recall that this is irrelevant on the Normic Support Account – what matters is whether it would be abnormal on the evidence, and that it ultimately isn't. Thus, we have good reasons for thinking that after learning D Helen's belief that the she will see her friend Bob next weekend is no longer normically supported by the evidence and subsequently fails to be justified. Another problematic result.

These two cases illustrate a *general structural problem* with the Normic Support Account – namely that negligible pieces of evidence can have serious defeating power. The heart of the Easy-Defeat Problem, to reiterate, is that according to the Normalcy View the introduction of a single normal E & not-P world is sufficient for defeat - no matter how remote the probability that it is the actual world. Additional cases illustrating this problem can easily be generated using the following recipe: First, chose a proposition P that is justified for S. Next, chose a proposition Q that is entailed by P. Finally, introduce a new piece of evidence D, such that (i) D provides an explanation for why Q might obtain and (ii) Q remains so overwhelmingly improbable that we judge D insufficient to defeat P's justificatory status.

5.2. *The Bitter Pill as a Global Version of the Easy-Defeat Problem*

So far, the Easy-Defeat Problem has been presented as a *localized* problem for certain isolated cases – e.g. Allergy and Lightning. In these cases, the problem has been local because the defeating evidence only affected a *particular* proposition relevant to the respective cases (e.g. propositions about peanut-contaminated flour or people being struck by lightning). But the Easy-Defeat Problem can of course also arise on a *global* level. On a global version of the Easy-Defeat Problem, a seemingly trivial piece of evidence will defeat the justificatory status of many (or even *all*) of the subject's beliefs, effectively leading to global skepticism about justification (and presumably knowledge). Under what conditions do these cases arise?

Global versions of the Easy-Defeat Problem arise when a piece of evidence establishes for a large number of propositions the existence of a very small number of normal worlds in which they are false. Put in terms of *explanations*, global versions of the Easy-Defeat Problem arise when a piece of evidence provides an explanation for why a large number of propositions may turn out to be false, even if it remains overwhelmingly improbable that any given belief is false. In such a case a seemingly negligible piece of evidence would result in an agent losing justification for a large number (or perhaps all) of their beliefs. Such cases, as we will see, make the Easy-Defeat Problem even more severe.

Can we provide an example for such a global version of the Easy-Defeat Problem? The pill scenario considered in chapter 3 provides such a case. Let's briefly reconsider the case

Pill S is given a bitter pill that ensures that a very small portion of S's ordinarily justified beliefs, let's say 1 out of every 10,000, chosen at random, will be false. The pill achieves this result by occasionally impairing S's cognitive connection to the evidence resulting in occurrences of, for instance, misperceptions or false memories. Importantly however these occurrences are incredibly rare. Finally, S knows about the effects of the bitter pill.

Recall that after taking the bitter pill Helen continues to form beliefs like everyone else and the vast majority of those beliefs, will of course turn out to be true. Helen will continue to form beliefs using the same belief-forming methods as before, i.e. perception, memory, testimony, etc. However, Helen knows that not all of her beliefs will be true. After all, the pill guarantees that a very small portion of them (0.001%) is going to be false. Importantly, Helen does not know which beliefs the pill will affect and since the effects of the pill are global it could be any of her beliefs.¹⁰⁵

¹⁰⁵ One way in which the pill might achieve its epistemic effect is by manipulating our cognitive connection to the evidence for the *one* belief out of 10,000 that will be false. In this case one can conceive of the pill's effects as manifesting through occurrences of misperceptions, false memories, etc. (See chapter 3 for a more thorough discussion of the pill scenario).

What result would the Normic Support Account of justification deliver in pill scenarios? On Smith's account the guarantee that 0.001% of beliefs will be false is sufficient to defeat justification for *any* belief Helen forms after having taken the pill. Put differently, on the Normic Support Account, knowing that one out of every ten thousand beliefs, chosen at random, is going to be false is sufficient to lead to global skepticism about justification (and presumably knowledge). It is not difficult to see why the account delivers this result: since S knows that for any post-pill belief there exists the possibility that it may have been affected by the pill, the subject's evidence is such that for any belief there exists a normal world in which the belief is false. This of course is the world in which that belief happened to be affected by the pill. Put in terms of explanations, since S is aware of the effects of the pill, S's body of evidence contains a possible explanation for why any of her beliefs may turn out to be false. As a result, none of the post-pill beliefs will be normically supported by the evidence and subsequently they will fail to be justified on the Normic Support Account.

This means that on the Normic Support Account the guarantee that just 1 belief out of 10,000, chosen at random, is going to be false, inevitably leads to global skepticism. This may strike us an excessively costly response to the pill scenario. Importantly then, the Easy-Defeat Problem can take the form of a *local* as well as a *global* problem.

5.3. *The Easy-Defeat Problem's Implications*

Why is the Easy-Defeat Problem bad news for the Normic Support Account? First, as the cases above illustrate, the problem of easy defeat shows that the account has counterintuitive implications. For instance, in Allergy we would not ordinarily consider the fact that 1 bag of contaminated flour has accidentally been put into circulation sufficient to defeat Helen's justification for believing that her brownie is safe to eat. We might accept that in light of such evidence Helen should *slightly* lower her confidence in P, but we would not accept that the evidence is anywhere near sufficient for defeat. Similar considerations apply in Pill. We do not ordinarily think that one should lose justification for all of one's beliefs after finding out that 1 randomly chosen belief out of 10,000 is going to be false. More generally, we do not

tend to assign defeating powers to highly improbable error possibilities. A plausible explanation for our reluctance to accept the normalcy view's defeat predictions might be that we generally expect an undermining defeater to make the falsity of a belief sufficiently probable – merely establishing that it would not be entirely abnormal if the belief turned out to be false seems to set the bar for defeat problematically low. Thus, if the Normic Support Account is to be understood as a descriptive account, i.e. an account that adequately captures our ordinary concept of justification, then the account fails; for in at least some cases - those involving low probability defeaters - its predictions systematically diverge from our ordinary judgments. If on the other hand the Normic Support Account is supposed to be prescriptive, i.e. an account about how we *should* think about justification rather than how we *do* think about justification, then one would have to argue that the account's counterintuitive easy-defeat predictions are virtues rather than vices of the view. Making this conclusion palatable however will not be an easy task.

Alternatively, we can give a diagnosis of the problem in terms of Ichikawa's notion of *stinginess*. In recent work, Ichikawa (2014) made explicit what is often taken for granted, namely that one condition any plausible theory of justification needs to satisfy is that it is not 'too stingy'. According to Ichikawa, an account of justification is too stingy if, "it denies justificatory status to too many beliefs that are intuitively justified" (186). What the Easy-Defeat Problem demonstrates is that by predicting defeat too readily the Normic Support Account systematically violates this plausible stinginess constraint. The case of Pill makes that issue particularly vivid.

To make matters worse, there are good reasons for thinking that instances of the Easy-Defeat Problem, like Lightning or Allergy, are not far-fetched isolated cases, but that we find ourselves in these kinds of epistemic situations frequently. This is the case because we frequently come across information that provides a possible explanation for why many of our beliefs may be false. Put differently, we frequently come across information that introduces, for a considerable number of beliefs, a small number of normal worlds in which the belief is false. Just consider the ubiquity of information like the following: 'Devastating fire in New York City apartment', 'Contaminated eggs discovered in supermarkets', 'Fatal hit and run in central London', 'Shark attack off the coast of Australia', 'Plane crash over the Atlantic, or

‘In very rare cases (1 in 100,000) this medication has been linked to anxiety and depression’. Since evidence of this kind provides an explanation for why one might falsely believe that one still has an apartment in New York City, that one’s friends who live in London are alive and well, that one’s cousin who lives in Australia and enjoys an occasional swim has not fallen victim to a shark attack, or that one’s allergy medication will make one feel better rather than worse, these beliefs would fail to be justified on the Normic Support Account. Moreover, one would also lose justification for any proposition entailed by these beliefs. So, the epistemic consequences of the Easy-Defeat Problem are much more expansive and far-reaching than initially thought. Of course, the falsity of any of these beliefs is highly improbable, but recall that the Normic Support Account is not sensitive to probabilistic considerations. Thus, the Normic Support Account combined with the fact that we frequently acquire evidence that provides an explanation for why some of our beliefs may be false entails that we lack justification for a great many beliefs that we ordinarily take to be justified.

The Easy-Defeat Problem then poses a two-pronged challenge against the Normic Support Account. In a first instance it gives rise to straightforward counterexamples to the view. And secondly, it reintroduces skeptical worries, which Smith was hoping to avoid.

5.4. Anticipating a Response:

A perhaps initially compelling response to the Easy-Defeat Problem might be to deny that in cases like Allergy, Lightning, or any of the other examples, there really does exist a normal world in which these beliefs are false – let’s call this the *Denial Strategy*. However, I think there are good reasons to be skeptical about this strategy. Recall that within the normalcy framework the notion of normality is explained via the notion of ‘calling for explanation’: for any proposition P and body of evidence E, there does not exist a normal world in which P is false, only if, given E it would call for an explanation if P were false. Since in Allergy, Lightning, Pill, as well as the other examples we considered, the evidence *does* contain an explanation for why the relevant belief that P may be false, it seems natural to assume that in these cases it

would no longer call for an explanation if P turned out to be false. The principle underlying this thought is the following.

Explanation If one's body of evidence E contains an explanation for why P may turn out to be false, then it would no longer call for an explanation if P turned out to be false.

So, if we accept Explanation, then in Allergy, Lightning, etc., there does exist a normal world in which P is false. This means that proponents of the *Denial Strategy* will need to reject Explanation. Explanation however seems very plausible and it is difficult to see on what grounds one might deny it.¹⁰⁶ Nevertheless, I want to briefly consider but reject one possible argument against Explanation.

¹⁰⁶ I thank an anonymous referee for *Philosophical Studies* for the interesting observation that Explanation may have the unintended, and perhaps problematic consequence, of not just explaining why, on the Normic Support Account, we lack justification in easy defeat cases like Lightning and Allergy but that it may also provide reasons for thinking that we lack justification in lottery-like cases discussed earlier.

The reviewer considers the following lottery-like case. Suppose I am justified in believing P, that I will cook dinner tonight. On the Normic Support Account this entails that I also need to be justified in believing Q, that there will not be a power outage tonight. Initially, we may think that this is unproblematic, because if Q turned out to be false and there was in fact a power outage tonight, then this would be abnormal or call for explanation. However, many of us who are justified in believing Q will also be justified in believing R, that sometimes trees fall on power lines and cause power outages. At this point, the referee suggests that one might think that R provides an explanation for why Q may turn out to be false. If this analysis is correct, then according to Explanation it would no longer call for a special explanation if Q was false and subsequently Q as well as P - since P entails Q - would fail to be justified on the Normic Support Account. This result would not just be bad news for the Normic Support Account, as lottery-like cases would present a more serious objection to the Normic Support Account than initially thought, but it would also mean that there is a tension between Explanation and the previous concession that the Normic Support Account can straightforwardly deal with the skeptical threat posed by lottery-like scenarios (Sect. 5.1). How can this tension be dissolved?

Even though Smith never gives a detailed analysis of the notion of *calling for an explanation* or what it is to *have an explanation for something* - that was Anderson's (2017) primary criticism of the Normic Support Account - it seems there is room to plausibly deny that R does in fact provide an explanation for why Q may be false. Here is how we may motivate this position in a way that I take to be in line with how Smith seems to think about these cases. In lottery-like cases - like the one suggested by the reviewer - the subject lacks any specific (or positive) evidence for thinking that in this instance their belief may actually be false. Put differently, in lottery-like cases the subject lacks any reason for thinking that the disobliging environmental conditions, which would make them falsely believe that P, may actually obtain. In this regard lottery-like cases differ from easy-defeat cases, in which D

Reconsider the Lightning case and now imagine Helen arrives, as planned, at the train station in Oxford where she is expecting Bob. After the platform has cleared Helen is surprised to find that Bob is not at the train station. Helen is now beginning to consider various reasons that would explain why Bob is not there to pick her up. She of course is aware that a man in Oxfordshire got struck by lightning, but of course there are other possible explanations, maybe Bob got stuck in traffic, or maybe he has simply forgotten Helen's exact arrival time. Considering these numerous alternatives Helen now really wants to know what happened – i.e. she wants to know which of the possible explanations for Bob not being there materialized. In short, despite having an explanation for why Bob may not be at the train station, one may think that for Helen, the situation still calls for an explanation. If this is the case, then proponents of the Normic Support Account may argue against Explanation as follows: as the considerations above show, at least sometimes it is possible for a body of evidence E to contain an explanation for why P may turn out to be false but, if P actually turned out to be false, this would still call for an explanation. So, Explanation is false.

This argument however is not as persuasive as it may initially seem. Recall that for Smith the 'calling for an explanation' relation is a metaphysical relation between a body of evidence and a proposition. It is not supposed to be a psychological notion that tracks whether our curiosity is satisfied or not. So, even if Helen experiences some residual curiosity about which of a number of candidate explanations actually explains why Bob did not show up at the train station, this is not the notion of 'calling for an explanation' that is relevant to the Normic Support Account. From the fact that Helen's curiosity is not satisfied it does not follow, that

provides a specific reason for thinking (or suspecting) that Helen's belief that P may be false, or that the disobliging environmental conditions that would make her falsely believe that P may actually obtain. Hence, it seems that we can reasonably deny that R, on its own, provides an explanation for why Q may be false.

Acknowledging Smith's (2016: 57) remark that his response to the skeptical threat posed by lottery-like cases is not as definitive as one might wish, I think that this is a *prima facie* plausible and charitable way of thinking about explanations that would (i) preserve the idea that the Normic Support Account is compatible with justification for lottery-like propositions and (ii) dissolve the alleged tension between Explanation and justification in lottery-like propositions.

her body of evidence, metaphysically speaking, does not contain an explanation for why Bob may not be at the station.

We can simplify this idea. Consider a subject whose body of evidence is such that it contains three possible explanations for why their belief that P may turn out to be false. Now consider the following way of reasoning: if P turned out to be false, then they would still want to know *precisely* why their belief that P turned out to be false; so, if P turned out to be false, it would still call for a special explanation. Hence, the belief that P is justified. However, this way of reasoning seems problematic, it would surely be wrong to conclude from the fact that someone's body of evidence contains three candidate explanations for why P may be false that their body of evidence does not contain an explanation for why not-P may obtain. This points back to the following more general thought: whether a body of evidence metaphysically speaking contains an explanation for why a proposition may be false comes apart from whether or not our curiosity would be satisfied in case the proposition is false. So, only because sometimes we would want to know more if a belief turned out to be false or because we would demand further explanation to satisfy our curiosity if a belief turned out to be false, it does not follow that our body of evidence, metaphysically speaking, does not already contain an explanation for why the belief may turn out to be false. Hence, once we recognize the difference between something *metaphysically requiring explanation* and something requiring explanation for personal reasons or in light of some psychological notion like our *curiosity*, we see why the above objection to Explanation fails. Thus, I don't think that there is a straightforward way out of the Easy-Defeat Problem.

In the next two sections I consider the second leading Normalcy View – viz. Leplin's *Normic Reliabilism* - and whether it can avoid the Easy-Defeat Problem. Ultimately, I conclude that it cannot.

6. *Normic Reliabilism*

Recall that according to Leplin, a belief is justified only if it was formed by a method that is perfectly reliable under normal conditions and one has no reason to believe that conditions are abnormal. To see what he has in mind, consider again some

paradigmatic instances of justified beliefs. If a belief based on perception were to be false, then, so the thought goes, under *normal conditions* perception would not have produced the belief - hence, the belief is *prima facie* justified. What does it mean for conditions to be normal according to Leplin?

Unlike Smith, who cashes out normalcy via a metaphysical relation (the ‘calling for an explanation’ relation) between a body of evidence and a proposition, Leplin explains normalcy in terms of the conditions that we naturally presuppose when using certain belief forming methods: “conditions normal for a method are conditions typical or characteristic of occasions and environments in which the method is usable or applicable” (37).

To clarify this idea, Leplin considers the helpful analogy of a thermostat. Ordinary thermostats, so the thought goes, are not reliable in all temperature ranges. Instead their reliability is limited to certain physical conditions (e.g. -40°C to $+50^{\circ}\text{C}$). In consulting a thermostat, we must presuppose that these conditions are met, for otherwise the thermostat will not be a reliable guide to an environment’s temperature. Thus, when consulting a thermostat, we must presuppose that we are in an appropriate physical environment for the thermostat to work properly. It follows that if we have reason to believe that any of the natural presuppositions ensuring that the thermostat is an inerrant guide to outside temperatures are violated - i.e. when we have reason to believe that we are in an abnormal environment where temperatures are below -40°C or above $+50^{\circ}\text{C}$ - then we can no longer assume that the thermostat is a perfectly reliable guide to outside temperatures. When this is the case, the thermostat becomes *unusable* for epistemic purposes (40).

This idea will of course easily generalize to other belief-forming methods. For instance, when we form beliefs based on perception we presuppose the absence of trick mirrors, barn facades, and deceptive lighting. Likewise, in cases of testimony we must presuppose the sincerity of the testifier, proper hearing, etc. (40). One could of course easily extend this list of natural presuppositions to include the absence of sudden power outages, barn facades, peanut-contaminated baking flour, dubious pills, etc. Importantly, it is these natural presuppositions which *ensure* that our belief-forming methods are perfectly reliable. Again, this explains why besides requiring a belief-forming method to be perfectly reliable under normal conditions, justification

also requires that the believer has no reason for believing that conditions are abnormal, i.e. that any of the natural presuppositions are violated; for in abnormal conditions a method's perfect reliability across modal space is no longer guaranteed.

In what follows I will argue that Normic Reliabilism, like the Normic Support Account, faces the problem of easy defeat.

7. *The Easy-Defeat Problem Strikes Again*

7.1. *Lottery Skepticism*

One unsurprising consequence of Normic Reliabilism is that like the Normic Support Account it denies the justificatory status of beliefs based on purely statistical evidence, e.g. lottery beliefs. After all, in the case of lotteries for instance one knows that the method by which one comes to believe that one's ticket is a loser is not perfectly reliable; in a world in which one happens to hold the winning ticket one would nevertheless believe that one's ticket is a loser. This once again raises the issue of Hawthorne/Vogel style lottery skepticism: If Normic Reliabilism denies justification in lottery propositions but many ordinary beliefs entail lottery propositions, then hardly any of our beliefs would be justified. However, like Smith, Leplin offers a response to the standard Hawthorne/Vogel skeptical challenge: Leplin, like Smith, denies the symmetry between *genuine* lottery-propositions and propositions that are merely *lottery-like*. While Smith distinguishes genuine lottery propositions from lottery-like propositions via the notion of 'calling for an explanation', for Leplin the relevant difference is that in genuine lottery cases it is guaranteed that the employed belief-forming method will produce a false belief, while in lottery-like case there is no such guarantee (2009: 104). So, while in genuine-lottery cases we know that there exists a possible world in which we falsely believe that our ticket is a loser, the belief-forming method that produce *lottery-like* beliefs (e.g. that Donald Trump has not recently suffered a heart attack) may well be perfectly inerrant across modal space. For this reason, we lack justification for believing genuine lottery proposition while beliefs in *lottery-like* proposition can be justified. So, traditional Hawthorne/Vogel lottery skepticism is avoided. However, as we will see, Normic Reliabilism too is

committed a notion of epistemic defeat on which seemingly irrelevant pieces of evidence are sufficient for defeat. Again, this gives rise to a new skeptical threat.

7.2. *The Easy-Defeat Problem Returns*

Recall that according to Normic Reliabilism, for a belief to be justified it is necessary that it is *reliably produced*. And a belief is reliably produced only if it is produced by a belief-forming method M that is perfectly reliable under normal conditions *and* the believer has no reason to believe that conditions are abnormal. It follows that the justificatory status of a belief is defeated as soon as the agent has reason for believing that conditions are abnormal. This is the case because under abnormal conditions, i.e. when any of the presuppositions that ensure a method's perfect reliability are violated, a method's perfectly reliability across modal space is no longer guaranteed. These considerations give rise to the following account of defeat for Normic Reliabilism.

Normic Reliabilism Defeat (NRD) S's justification for believing P is defeated if, given a new piece of evidence D, the believer has reason for believing that conditions are (or were) abnormal - i.e. if the believer has reason for believing D and D violates *any* of the presuppositions which ensure that the belief-forming method M which produced P is (or was) *perfectly* reliable across modal space.

Note the strength of the account. Since Normic Reliabilism demands nothing short of perfect reliability, having reason to believe that *any* of the reliability ensuring presuppositions has been violated is sufficient for defeat – even if the belief-forming method remains highly reliable. More precisely, justification is defeated according to (NRD) if one has reason to believe that there exists just a single world in which the method would produce or sustain a false belief. This gives rise to the same concerns previously raised for the Normic Support Account: Seeing that the account is entirely insensitive to probabilistic considerations, nothing prevents possibilities that are

overwhelmingly improbable from having serious defeating powers. This of course is just the *Easy-Defeat Problem* all over again.¹⁰⁷

To illustrate, consider what normic reliabilism would predict in response to the augmented lottery variants considered earlier. Let's start with the local cases. In Allergy, Helen formed the belief that the brownie is safe to eat based on the fact that it was clearly labelled 'peanut-free'. Let's assume that under normal conditions this belief forming method is perfectly reliable and one of the natural presuppositions that ensures Helen's method is reliable is something along the lines of, 'brownies labelled peanut-free will not be contaminated'.¹⁰⁸ However, in Allergy Helen has reason to believe that this presupposition has been violated or that conditions are abnormal: Knowing that there is a bag of contaminated flour in circulation gives Helen reason to believe that her belief-forming method is no longer *perfectly reliable* or inerrant across *modal space*. Thus, in Allergy Helen would lack justification for believing that her brownie is safe to eat. The same considerations apply *mutatis mutandis* to Lightning. In order for Helen's belief forming method (whatever it may be in this case) to be perfectly reliable, we must presuppose that people do not normally get struck by lightning. However, the evidence D, that a man in Oxfordshire has been fatally struck by lightning provides reason to believe that in this case the presupposition has been violated, i.e. that conditions are abnormal; for D provides reason to believe that Helen's belief forming method is no longer perfectly reliable across modal space. Helen will falsely believe that she will see her friend Bob next weekend in the world in which Bob happened to be the man in Oxfordshire who was struck by lightning. Hence, Helen's belief that she will see Bob next weekend fails to be justified on Leplin's account. The Easy-Defeat Problem has returned.

¹⁰⁷ For a brief discussion about Leplin's requirement of perfect reliability perhaps being too strict, see Christensen (2007).

¹⁰⁸ Leplin acknowledges that it is difficult to specify precisely the conditions under which a certain belief forming method is perfectly inerrant (41). This difficulty becomes apparent here, for it is not at all clear how the content of the relevant presupposition should be formulated. Considering that these presuppositions play a crucial role in Leplin's theory, we might worry that without a method for specifying their content his account remains somewhat incomplete. However, for current purposes I will ignore any difficulties surrounding this issue.

What about the global version of the Easy-Defeat problem? Do they also cause trouble for Normic Reliabilism? It appears they do. Recall that in Pill the subject's epistemic situation is such that 1 belief out of every 10,000 is guaranteed to be false. Importantly, the pill's effects are *global* and *random*, so any belief, regardless of the method that produced it, could possibly be affected by the pill. Now, as mentioned above, in order for our belief-forming methods to be perfectly reliable across modal space we need to presuppose the absence of such drugs, which threaten to undermine the perfect reliability of our belief forming methods. After all it is only under normal conditions that our belief-forming methods are perfectly reliable. But, since the subject knows about the effects of the pill, they have reason to believe that their belief-forming methods are no longer perfectly reliable across modal space. Hence, none of the subject's beliefs formed after taking the pill (bracketing those used in forming beliefs in necessary truths) will be justified according to Normic Reliabilism. Recall that this is the same unpalatable result that the Normic Support Account delivered. Normic Reliabilism then is equally liable to global versions of the Easy-Defeat Problem.

Recall that the problem of easy defeat presents a two-pronged challenge. First it exposes Normic Reliabilism to straightforward counterexamples in which the justificatory status of beliefs we ordinarily take to be justified is denied. And second it opens the door to a skeptical threat: Normic Reliabilism combined with the fact that we frequently acquire reasons for believing that our belief-forming methods are not perfectly reliable across modal space, yields the unpalatable conclusion that justified beliefs are a lot scarcer than we ordinarily think.

As we have seen, the two leading Normalcy Views - the Normic Support Account and Normic Reliabilism - both face the problem of easy defeat. On both accounts insignificant pieces of evidence turn out to have considerable defeating powers. This makes the accounts highly stingy, leads to counterexamples, and exposes them to skeptical threats, which both Smith and Leplin were hoping to avoid. Is there a way around the Easy-Defeat Problem? In the next section I propose modifications to the two accounts that would allow them to avoid the Easy-Defeat Problem. However, we will see that the required modifications will fail to save the accounts.

8. *Can the Easy-Defeat Problem be Avoided?*

8.1. *The Easy-Defeat Dilemma*

What got the Normic Support Account into trouble is that in order for a belief that P to be justified, P must be true in *all* normal E worlds. As far as I can see, the only way to avoid the problem of easy defeat is to *weaken* the notion of normic support so that it can accommodate a small number of normal worlds in which P is false. A straightforward way of achieving this goal would be to concede that in order to be normically supported by the evidence a belief that P must be true not in *all* but only in *nearly all* normal E worlds. Weakening normic support's modal strength in this way yields the following *Weak Normic Support Account*.

Weak Normic Support Account (WNSA) In order for one to have justification for believing a proposition P, it is necessary that one's body of evidence E normically support P – it is necessary that **nearly all** of the most normal worlds in which E is true are worlds in which P is true.

Similarly, what gets Normic Reliabilism into the problem of easy defeat is that in order for a belief to be reliably produced it is necessary that (i) the belief-forming method is perfectly reliable under normal conditions, and that (ii) the believer has no reason to believe that conditions are abnormal. As we have seen (ii) can easily be violated by seemingly insignificant pieces of evidence because according to Normic Reliabilism one has reason to believe that conditions are abnormal as soon as one has reason to believe that any of the presuppositions that ensure a belief forming method's perfect reliability across modal space have been violated. To avoid this problem, we can weaken (ii) so that if the believer has reason to believe that conditions are abnormal justification can be retained if the believer also has reason to believe that the relevant belief-forming method remains *highly reliable*. This weakened version of Normic Reliabilism can be expressed as follows.

Weak Normic Reliabilism (WNR) In order for S to have justification for believing a proposition P, it is necessary that S's belief was reliably produced – it is necessary that (i) the belief that P was produced by a belief-forming method M that would not, in any normal world, have produced the belief that P if P were false and (ii) that the believer has no reason to believe that conditions are abnormal **or** if the believer has reason to believe that conditions are abnormal, it is reasonable to assume that the method remains *highly reliable* across modal space.

These weakened versions of the Normic Support Account and Normic Reliabilism can avoid the Easy-Defeat Problem. Consider for instance what (WNSA) and (WNR) would predict in Allergy. Despite the fact that Helen's evidence is compatible with a few normal worlds in which her brownie is contaminated and therefore not safe to eat, it will nevertheless remain the case that in *nearly all* E-worlds P is true. So, according to (WNSA) Helen's belief would be justified. Similar considerations apply to (WNR). Even though Helen has reason to believe that conditions are abnormal – after all one of the presuppositions that ensure her belief-forming method's perfect reliability across modal space has been violated – she also has reason to believe that her belief-forming method will remain *highly reliable* as the number of worlds in which she falsely believes that her brownie is safe to eat is very small. Analogous considerations apply in the case of Lightning. Given Helen's total evidence, her belief that she will see Bob next weekend, while not true in *all* normal worlds, will nevertheless be true in *nearly all* normal worlds. So, according to (WNSA) Helen would be justified in believing that she will see Bob next weekend. Again, similar considerations apply in the case of (WNR). Even though Helen has reason to believe that one of the presuppositions ensuring her belief-forming method's perfect reliability has been violated, she also has reason to believe that the method by which she came to believe that she will see Bob next weekend will remain *highly reliable* across modal space. So, according to (WNR) her belief will be justified. Thus, the weakened versions of the two accounts appear to deliver the intuitively correct results and avoid the Easy-Defeat Problem. However, while initially promising, there are good reasons to be suspicious about the weakened accounts.

(WNSA) and (WNR) essentially mark a return to a Lockean View of justification. An indication of this collapse into threshold views is the ‘*nearly all*’ quantifier in (WNSA) and the appearance of ‘*highly reliably*’ in (WNR). One consequence of this is that the weakened accounts are no longer instances of the Normalcy View - they betray (NV). A related second consequence is that like other Lockean views, the weakened accounts face the problem of *risk accumulation* over conjunctions. This means that (WNSA) and (WNR), unlike their stronger counterparts, will not be compatible with (MPC). This is problematic because it was the promise of preserving (MPC) that was the primary motivation for the two accounts. Without preserving (MPC) the accounts become unmotivated.

It appears then that the Normic Support Account and Normic Reliabilism face a dilemma: The original accounts, which are committed to (NV), face the Easy-Defeat Problem, while the weaker versions, which by violating (NV) are able to avoid this problem, collapse back into threshold views and therefore cannot make good on their motivating promise of preserving (MPC). We may call this the *Easy-Defeat Dilemma*.

8.2. Generalizing the Problem

Besides putting pressure on the Normic Support Account and Normic Reliabilism, the above considerations also point towards a more general connection between normalcy views and the Easy-Defeat Problem. As we have seen, what gets the two leading normalcy views into trouble has little to do with any *account-specific-features* (e.g. their respective notions of normality) and everything to do with a more *general feature* of the views, namely their modal strength. Importantly, the modal strength of normalcy views is not merely an optional feature of the views but instead it is required to make good on the promise of preserving (MPC) – after all the problem of risk accumulation needs to be avoided. In other words, the modal strength of (NV), which is directly responsible for the Easy-Defeat Problem, is an *essential feature* of Normalcy Views of justification. As a result, we can expect the problem to generalize to any Normalcy View – i.e. to any account of justification committed to (NV) - regardless of how the details are developed. What does this mean for the general prospects of Normalcy Views? Well, in so far as we take seriously the idea that

justification should be sufficiently robust and be able to survive small amounts of risk, we have good reasons to be skeptical about the success of Normalcy Views.

Can the Easy-Defeat Problem be generalized even further? And if so how far is its reach? I briefly want to outline what I take to be good reasons for thinking that the Easy-Defeat Problem will generalize to *any* view of justification committed to (MPC). If this is correct, then we should view the Easy-Defeat Problem as an inherent cost of preserving (MPC).

Recall that in order to preserve (MPC) a theory of justification needs to avoid the problem of risk accumulation. In other words, in order to preserve (MPC) a theory of justification needs to avoid counting beliefs that carry with them any relevant¹⁰⁹ degree of risk as justified. Put in different terms, theories of justification that want to preserve (MPC) will need to deny justification as soon as risk – regardless of its degree – enters the picture. If a theory of justification is made compatible with *any* degree of relevant error risk, (MPC) will fail. And this is precisely why we can expect the Easy-Defeat Problem to generalize to any theory of justification compatible with (MPC). If we want to preserve (MPC) we cannot make justification compatible with risk, and if we cannot make justification compatible with risk, then we will need to accept that justification will be lost very easily. It is difficult to see how any theory of justification that claims to be compatible with (MPC) will be able to avoid the problem of denying justification in cases like Allergy, Lightning, and Pill. Hence, we can expect the Easy-Defeat Problem and the cases we used to illustrate to generalize widely to any theory of justification committed to (MPC).

9. Is 'Being Normally True' Necessary for Justification?

¹⁰⁹ I use the term 'relevant risk' to capture whatever risk amounts to on different accounts of justification. For instance, on the Lockean View risk will be spelled out probabilistically – i.e. on the Lockean View a belief has a relevant degree of risk associated with it if it has a non-zero probability of being false. In contrast, on the Normalcy View, risk will be spelled out in terms of normal worlds in which one falsely believes that P. However, for our current purposes we can abstract away from these issues and just talk about 'relevant risk' to pick out the notion of risk relevant to any given account we may be interested in.

In this chapter we have looked at the prospects of non-probabilistic alternatives to the Lockean View, which promise to preserve (MPC). The perhaps most promising candidates appear to be instances of what we called the Normalcy View. By parting with the orthodox approach of explaining justification in terms of high probability and by identifying a different justification-conferring property (J-property) - the property of *being normally true* – normalcy views avoid the problem of risk agglomeration, and, as a result, unlike the Lockean View, are able to preserve (MPC). Since the preservation of (MPC) is a plausible desideratum for a satisfactory theory of justification normalcy views look initially very attractive. In fact, the Normalcy View may well be the most promising framework for thinking about justification in a way that is compatible with deductive closure to date.

However, as we have seen, the normalcy picture of justification faces a serious objection. I argued that normalcy views of justification - i.e. accounts of justification committed to (NV) - face a serious obstacle when it comes to the notion of epistemic defeat. By requiring that in order for a belief to be justified there cannot exist a single normal world in which the belief is false, normalcy views become highly sensitive to defeating evidence. And as a result, justification is lost much too easily. This, I argued, exposes normalcy views not just to straightforward counterexamples but to an uncomfortable skeptical threat. In the end, it appears that the Normalcy View, despite being initially attractive is simply too demanding. To borrow Ichikawa's terminology again, what the problem of easy defeat illustrates is that on the Normalcy View, justification becomes too stingy – i.e. it denies the justificatory of too many beliefs that we ordinarily take to be justified. Thus, we have good reason to reject the Normalcy View's central claim that in order for a belief to be justified there must exist no normal worlds in which the belief is false. Put in terms of properties, we should resist the claim that in order for a belief to be justified it is necessary that it has the property of *being true in all normal world*.

Chapter 5

Exploring New Territory: Pluralism about Justification and The Functional Theory

1. *Probabilistic vs Non-Probabilistic Accounts to Justification: A Dilemma for Monists*

In the previous chapters we looked at two very different ways of answering the question ‘what justifies belief?’ but found reasons to be unhappy with both.

According to the most prominent accounts of epistemic justification, the property that makes a belief justified or the *justification-conferring* property – henceforth the *J-property* - is some property along the lines of *being highly probable for S*. We called this probabilistic picture of justification the Lockean View.

Lockean View A belief that P is justified for S iff P is highly probable for S.

While the Lockean View has a lot of intuitive appeal, we have seen that it is not without problems. The perhaps most serious objections against explicating justification in probabilistic terms is that the resulting theories will be incompatible with multi premise closure – a principle that we have argued is a plausible desideratum for a satisfactory theory of justification.

Multi Premise Closure (MPC) If S is justified in believing p and S is justified in believing q... and S is justified in believing n, then S is justified in believing the conjunction (p & q... n).

The main reason for thinking that the Lockean View is incompatible with (MPC) is the problem of risk agglomeration: as we begin conjoining individual beliefs, the error-risk associated with each belief starts to accumulate and eventually – after conjoining sufficiently many belief – we end up with a conjunction that is highly improbable and therefore will fail to be justified on the Lockean View. Thus, any account that explains justification in terms of high probability will not be able to accommodate (MPC).

So, if we want to preserve (MPC), then the J-property needs to be something other than the property of *being highly probable*. More specifically, in order to preserve (MPC) an account of justification needs to avoid the problem of risk accumulation. We have considered two families of views that fit this mould and that offer the promise of preserving (MPC).

The first one was the *Strong View*, which we quickly dismissed because it failed to satisfy the Fallibility Platitude. The second, and considerably more promising one, was the *Normalcy View*, versions of which have recently been defended by Jarrett Leplin (2009) and Martin Smith (2010, 2016). According to these views, broadly speaking, a belief that P is justified for S only if, given S's evidence, there exists no normal world in which S falsely believes that P.

Normalcy View The belief that P is justified for S only if, given S's evidence E, there does not exist a single *normal world* in which S falsely believes that P.

For proponents of the Normalcy View, the property that justifies beliefs – the J-property - is the property of *being true in all normal worlds*; or, for short, of *being normally true*.

An important feature of the Normalcy View is that it explains justification in non-probabilistic terms. In order to be justified, a belief that P must not just be true in *most normal worlds*, it needs to be true in *all* normal worlds. As a result, the Normalcy View avoids the problem of risk accumulation. Since in order for a belief that P to be justified there must not exist a single normal world in which S falsely believes that P, there will not exist a single normal world in which a conjunction of two (or more) individually justified belief will be false; after all, this would require a normal world

in which one of the conjuncts is false and this is ruled out by the Normalcy View. Hence, the Normalcy View, unlike the Lockean View, can preserve (MPC).

However, as we have seen in the previous chapter, normalcy views, due to the modal strength required to make good on the promise of preserving (MPC), face a serious objection: a belief fails to be justified if there exists just a single normal world in which S, given the evidence, falsely believes that P – even if the falsity of P is overwhelmingly improbable. This, we argued, makes normalcy views highly sensitive to defeating evidence and leads them to deny justification for many beliefs we ordinarily take to be justified. In short, accounts following the Normalcy View will be too stingy.

To get an initial feel for this worry we can consider a subject who holds a single ticket in a large lottery. Many have the intuition that the subject can justifiably believe that they will not win the lottery – let's call this a *lottery belief*. While Lockean views of justification have no problem accommodating the justificatory status of lottery beliefs – after all their truth is highly probable – normalcy views will not be able to deliver this result. In the case of a lottery, there does exist a normal world in which one falsely believes that one will lose; this of course is the world in which one happens to hold the winning ticket. So, regardless of the size of the lottery, accounts committed to something like the Normalcy View will end up denying that one can justifiably believe that one will lose the lottery. This problem of course generalizes to any belief supported by purely statistical evidence. Beliefs based on statistical evidence, no matter how strong the statistical evidence, will never be justified on the Normalcy View.

To make matters worse, this is not just an isolated problem for lotteries and cases involving naked statistical evidence. But instead, as we have seen in the previous chapter, the problem is more general in nature. On the Normalcy View, one loses justification for believing that P as soon as there exists just a single normal world in which S falsely believes P – even if one's belief remains overwhelmingly probable. One consequence of this feature of the Normalcy View is that seemingly trivial pieces of evidence can turn out to have considerable defeating powers. We called this the *Easy-Defeat Problem* and provided a number of cases (e.g. Allergy and Lightning) to illustrate the problem. What the Easy-Defeat Problem shows, is that normalcy views

will be overly stingy, as they deny the justificatory status of many beliefs that we ordinarily take to be justified.

At this point then it appears we face a dilemma in formulating a satisfactory theory of justification. Let's call it the (MPC)-Stinginess Dilemma.

(MPC)-Stinginess Dilemma If we go down the probabilistic route and opt for a Lockean account – i.e. we hold that the J-property is some property along the lines of *being highly probable for S* - then we will not be able to preserve (MPC). Alternatively, if we go down the non-probabilistic route and opt for something like a Normalcy View – i.e. we hold that the J-property is some property along the lines of *being true in all normal worlds* - then justification becomes too demanding and a lot of beliefs we ordinarily take to be justified will turn out to be unjustified.

I take this dilemma to present one of the central obstacles we face in providing a satisfactory theory of justification. How should we respond to the (MPC)-Stinginess Dilemma? The remainder of this chapter sets out to elaborate and defend a novel response to this question.

2. *The Eliminativist Strategy*

One strategy for responding to the (MPC)-Stinginess Dilemma would be to simply bite the bullet and to accept one of the horns of the dilemma. There are two options.

Bite-the-Bullet-1 Accept a version of the Lockean View and deny (MPC).

Bite-the-Bullet-2 Accept a version of the Normalcy View and make do with a stingy account of justification.

However, I think that there are reasons to be unhappy with both of these options. Why? Both of them force us to give up plausible desiderata for a wholly satisfactory theory of justification. Bite-the-Bullet-1 forces us to give up (MPC). And Bite-the-

Bullet-2 forces us to accept a theory of justification that is too stingy. So, both Bite-the-Bullet-1 and Bite-the-Bullet-2 should strike us as sub-optimal.

Now, even if one thinks that we could ultimately come to terms with the denial of (MPC) or with a theory of justification that is very demanding, it is easy to appreciate the fact that we could conceivably do better. A more satisfactory solution would be one on which justification would be somehow (i) compatible with a version of (MPC) without (ii) being too stingy.

The remainder of this chapter explores a different solution to the (MPC)-Stinginess Dilemma; one that promises to satisfy both (i) and (ii).

3. *The Pluralist Strategy*

One implicit, yet crucially important, assumption underlying the (MPC)-Stinginess Dilemma is that in providing a theory of justification we must choose between the Lockean picture and the Normalcy picture and that we cannot have a probabilistic and a non-probabilistic notion of justification working side-by-side. In other words, the dilemma presupposes that there is only one way for a belief to be justified or, more generally, that there exists only one justification-conferring property – i.e. one J-property. Let's call this assumed monism about justification, *J-monism*.

J-monism There is only one way for a belief to be justified – i.e. there is only one property that can make a belief justified

While J-monism is usually taken as the default position, it is ultimately optional and can reasonably be rejected. A rejection of J-monism yields pluralism about justification (henceforth, *J-pluralism*).

J-pluralism There is more than one way for a belief to be justified – i.e. there is more than one property that can make a belief justified.

This pluralist picture of justification offers interesting new resources for responding to the (MPC)-Stinginess Dilemma. The key idea is the following: on a pluralist picture

of justification we can accommodate the notion that some beliefs are justified in virtue of *being highly probable* whilst other beliefs are justified in virtue of *being normally true*. How does this help solve the dilemma?

Note that for all beliefs that are justified by having the property of *being normally true* justification will be closed under multi premise deduction – after all this was the primary motivation for normalcy views of justification. This means that for all beliefs that would be justified on normalcy views - e.g. ordinary instances of beliefs based on perception, testimony, and memory - (MPC) will hold. Importantly however, we now no longer need to accept the problematic consequence that any belief that fails to be true in all normal worlds will necessarily be unjustified. Beliefs that are not true in all normal worlds may still be highly probable for S – they may be true in *nearly all* normal worlds - and therefore can still be justified.¹¹⁰ So, by making room for the idea that there can be more than one way for a belief to be justified, J-pluralism offers the promise of providing an attractive new way out of the dilemma - we can preserve a version of (MPC)¹¹¹ whilst avoiding the issue of justification becoming too stingy. As a result, J-pluralism can overcome one of the central challenges for J-monism. On a pluralist picture of justification, it seems that we can have our cake and eat it too.

In what follows I will make the pluralist proposal more precise. An important question that still needs to be answered is what the cooperation between the two different ways of being justified looks like. In other words, we still need to answer the question of which beliefs are justified by which property.

3.1 An Analogy

To answer these questions, it will be helpful to briefly consider a previous, well-known, pluralist project: *Alethic Pluralism* or *Truth Pluralism* championed by Crispin

¹¹⁰ The same applies the other way around. A belief may be highly unlikely and therefore lack the property of being highly probable for S; but, it may still have the property of being true in all normal worlds and in this case would be justified.

¹¹¹ We will of course need to localize (MPC) as it will only hold for beliefs justified in a certain way, i.e. for beliefs that are justified in virtue of *being normally true*. We will turn to this issue in Sect 5.4.

Wright (1992; 1999) and Michael Lynch (2004; 2009). Since truth pluralism and J-pluralism have similar motivations the analogy will be helpful in providing an initial blue print for how to get a pluralist theory off the ground.

The key motivation for truth pluralism – the idea that there is more than one way for a belief to be true - is the apparent diversity of content amongst the propositions we believe and take to be true: snow is white, murder is wrong, $2+2=4$, etc. Each of these sentences appears to be true, yet it seems implausible that they should all be true in the same way. Put differently, it seems implausible that all of the propositions are made true by one and the same property. For instance, the sentence ‘snow is white’ seems to be made true by its accurately *representing* some feature of the external world. However, it seems less convincing to try and explain the truth of moral or mathematical statements by their accurate representation of a mind-independent reality. Instead their truth is often explained by appeal to a different kind of property - usually some form of *coherence* with a larger belief-system.¹¹² Importantly, neither representational theories of truth, nor coherence theories of truth alone strike us as wholly satisfactory – Lynch (2009) calls this the *Scope Problem*. These considerations have prompted truth pluralists to reject alethic monism, the idea that there is only one way in which a proposition or a belief can be true, and moved them towards the idea that there is more than one way for a belief to be true or that there is more than one property that can make a proposition (or a belief) true.

According to the alethic pluralist, some beliefs are true in virtue of accurately representing the external world, while other beliefs are true in virtue of their coherence with a larger belief system. More specifically, truth pluralists argue that the property that makes a belief or a proposition true depends on the *domain of inquiry*. In some domains (e.g. beliefs about the external world) truth is realized by a kind of accurate representation of the external-world, while in other domains (e.g. beliefs about morality or mathematics) truth is realized by some form of coherence (Lynch 2009:

¹¹² For instance, Lynch (2009: 39) proposes the following version of a coherence theory of truth

(SC): The belief that p is true if and only if that belief is supercoherent

Where, “a belief is supercoherent just when it is a member of a coherent system of beliefs at some stage of inquiry which would remain coherent without defeat in every successive stage of inquiry. (2009: 40).

5). By making room for the idea that depending on the domain of inquiry beliefs can be true in different ways, truth pluralists are able to accommodate the apparent diversity amongst the proposition we take to be true.

As we can see, alethic pluralism and J-pluralism share many important features in common. Both start out by acknowledging that there are at least two competing theories, which both capture something important about the nature of truth. For alethic pluralists these are representational and coherence theories of truth; for the J-pluralism that I am concerned with here, they are probabilistic and non-probabilistic accounts of justification. However, neither theory is wholly satisfactory, as neither of them captures all the desiderata we want to capture. For alethic pluralists neither representational nor coherence theories of truth can account for all the propositions we take to be true; for J-pluralists neither probabilistic nor non-probabilistic accounts can account for all the belief we take to be justified – probabilistic theories give out in cases of long conjunctions, and the leading non-probabilistic theories which promise to avoid this problem give out any time they face any degree of relevant risk. Truth pluralists call this the *Scope Problem*, we called the justification analogue of the Scope Problem the *(MPC)-Stinginess Dilemma*.

Note that the two problems are similar in nature. We could easily present the scope problem as a dilemma: if we opt for a representational theory of truth, then we will have to deny that any propositions with non-representational content – e.g. propositions about morality or mathematics – can be true; alternatively, if we opt for a coherence theory of truth, then the truth of all propositions would have to be explained in non-representational terms and this seems to get things wrong for large number of beliefs (e.g. beliefs about the external world). So, the Scope Problem can be formulated as dilemma closely analogous to the (MPC)-Stinginess Dilemma that motivates J-pluralism.¹¹³

¹¹³ Alternatively, we could follow Pedersen (2013) and present the (MPC)-Stinginess Dilemma in the form of a scope problem. Here is Pederson's general formulation of the epistemic scope problem.

Epistemic Scope Problem: epistemic value bearers x_1, \dots, x_n should all be classified as enjoying positive epistemic standing F , but no single theory of F can plausibly account for all of x_1, \dots, x_n being F . (2013: 52)

In the case we are interested in the epistemic value bearers would be beliefs (in both atomic propositions and conjunctions) and the positive epistemic standing F that they should all enjoy is *being justified*.

The force of the Scope Problem ultimately motivated truth pluralist to give up on the monist idea that there is only one way for a belief to be true and moved them towards pluralism. In Lynch's words, "In recent decades, many philosophers have come to think that the monist's quest for the nature of truth is a fool's errand." (2009: 4) Similarly, J-pluralism is motivated by the force of the (MPC)-Stinginess Dilemma outlined above. What the dilemma shows is that it is unlikely that we will be able to accommodate all the plausible desiderata for a satisfactory theory of justification on a monist theory. Echoing Lynch, proponents of J-pluralism may suggest that the monist's quest for the nature of justification too is a fool's errand. Alethic pluralism and J-pluralism then, may rightfully be seen as kindred spirits.

How does the alethic pluralist suggest the two notions of truth interact? In other words, how does the alethic pluralist suggest we decide, for any given proposition, what the relevant truth property is that determines whether the proposition is true?

To answer this question, the truth pluralist introduces the notion of *domains of inquiry*. Sometimes we are interested in the external world, other times we are interested in morality, and yet other times we are interested in mathematics; according to the alethic pluralist, these are simply different domains of inquiry. And what makes a belief true can vary depending on the relevant domain of inquiry. How exactly are these domains individuated? While this is a somewhat delicate matter for alethic pluralist, the standard answer is usually that domains of inquiry are individuated by the *content* of the propositions that fall within them. In other words, domains of inquiry are individuated, roughly speaking, by what its member propositions are *about*. For instance, when a proposition is about the external world, then, its truth value is determined by whether it accurately represents that world, and if a proposition is about morality or mathematics, then its truth value is determined by whether or not it coheres with a larger belief system.

Can pluralists about justification offer a similar story about which of the two justification-conferring properties is operative when? I argue that they can.

But as we have seen, no single theory can deliver the result that all beliefs that should be justified will actually be justified. Theories following the Lockean View cannot account for the justification of long conjunctions and theories following the Normalcy View cannot account for the justification of lottery beliefs and beliefs in Easy-Defeat cases.

3.2 Epistemic Environments and the Refined Proposal

While I argued in the previous chapter that we have good reasons to reject the central claim of the Normalcy View, namely that in order for a belief to be justified there cannot exist a single normal world in which, given the evidence, one falsely believes that P - it nevertheless seems that Smith and Leplin were on to something interesting when they observed that sometimes our body of evidence is such that there exists no normal world in which we falsely believe that P and other times our evidence is such that there do exist a number of normal worlds in which we falsely believe that P. One of the key insights that Smith and Leplin offer, we might say, is that just as there exists diversity amongst the *domains of inquiry*, so too does there exist a diversity of *epistemic environments*. With respect to some propositions we are in an epistemic environment in which, on the evidence, there exists no normal world in which we falsely believe that proposition – let's call this a *normal epistemic environment*.

Normal Epistemic Environment For any proposition P, an agent is in a normal epistemic environment with regards to P iff, on the evidence, there does not exist a single normal world in which P is false.

For other propositions however, we are in an epistemic environment in which, given one's evidence, there does exist at least one normal world in which we falsely believe that P - let's call this a *risk accumulating epistemic environment*.

Risk Accumulating Epistemic Environment For any proposition P, an agent is in a risk accumulating epistemic environment with regards to P iff, on the evidence, there does exist at least one normal world in which P is false.

So, like the truth pluralist who distinguishes between different domains of inquiry, the J-pluralist can distinguish between epistemic environments. How are epistemic environments individuated?

Epistemic environments, as mentioned above, are individuated by the total body of evidence: if, given the evidence E, there exists no normal world in which S falsely believes that P, then one is in a normal epistemic environment, otherwise one is in a risk accumulating epistemic environment. Since we have defined our epistemic environments *exhaustively*, it follows that for every proposition we are either in a normal epistemic environment or in a risk accumulating one,¹¹⁴ but never both. More succinctly, for any body of evidence E and proposition P, we will always be in *one and only one* epistemic environment with regards to P.

With the notion of epistemic environments – essentially an epistemic analogue of the alethic pluralists’ *domains of inquiry* - in place, we can now answer the question of how the two J-properties – the property of *being normally true* and the property of *being highly probable* – cooperate. In a normal epistemic environment, the property that makes a belief justified is the property of *being normally true*; and in a risk accumulating epistemic environment the property that makes a belief justified is the property of *being highly probable*. Let’s briefly consider how this more refined picture of J-pluralism deals with some of the cases considered earlier.

As proponents of the Normalcy View have argued compellingly, for a great many of our ordinary beliefs – e.g. beliefs based on perception, memory, and testimony - there usually does not exist a normal world in which we falsely believe that P. Hence, for these kinds of beliefs we are usually in a *normal epistemic environment*. According to the pluralism I propose, the property that makes these beliefs justified is the property of *being normally true*. Importantly, since the property of *being normally true* satisfies (MPC), all justified beliefs formed in normal epistemic environments will be closed under multi premise deduction.

However, we are not always this fortunate when it comes to our epistemic situation. Recall that for beliefs based on statistical evidence (e.g. lottery beliefs) and the relevant beliefs in easy-defeat cases - e.g. Lightning, Allergy, or Pill – our

¹¹⁴ There may of course be difficult cases where we are unsure whether or not there exists a normal world in which a given belief is false. Smith (2016) considers the example of a high ranked tennis player losing to a significantly lower ranked tennis player. Is there a normal world in which this happens? Who knows. But this is a different kind of issue. We may think of this problem as a normalcy analogue of the threshold problem for Lockean views of justification. For now, I think we should not make too much of these difficulties.

evidence is such that there does exist a small number of normal worlds in which we falsely believe that P. For these propositions we are not in a normal epistemic environment; instead we are in a risk accumulating environment. And in a risk accumulating environment the property that makes a proposition justified is the property of being *highly probable*. Since in the case of lottery beliefs and in easy-defeat cases there only exists a very small number of normal worlds in which one falsely believes that P, these beliefs remain highly probable for S. And as a result, lottery belief as well as the relevant beliefs in easy-defeat cases will be justified. However, due to the problem of risk agglomeration the property of *being highly probable* does not satisfy (MPC). It follows that multi premise deduction will not hold as a general principle for beliefs formed in a Risk Accumulating Epistemic Environment.

As we can see, what gets monists about justification into trouble - what forces them into the (MPC)-Stinginess Dilemma - is that they expect one justification-conferring property to apply across the board in all epistemic environments. In contrast, a key feature of the pluralism proposed here and the feature that ultimately allows it to avoid the dilemma is that it recognizes and takes seriously the previously ignored diversity of epistemic environments.

4. *Problems for Simple Pluralism*

At this point someone might think that J-pluralism should be understood as follows: there are two different *kinds*¹¹⁵ of justification – justification_n, according to which a belief is justified iff it is highly probable, and justification_p, according to which a belief is justified only if it is *normally true*. Let's call this view *Simple J-Pluralism*.

Simple J-Pluralism (SJP) There are simply two different kinds of justification: justification_n and justification_p

¹¹⁵ This could be understood either in terms of distinct properties or concepts.

Following (SJP) the term ‘justification’ would ambiguously pick out different kinds of justification.¹¹⁶ However, from the truth pluralism debate we know that if alethic pluralism is construed to simply mean that there exist a number of different truth predicates – e.g. truth_1 and truth_2 – then a number of issues arise. The most prominent ones are the *mixed inference problem* and the *problem of mixed compounds*. Since these problems are quite general in nature we may reasonably expect them to apply equally to (SJP). In this section I will consider these objections and see whether they really do pose a problem for (SJP).

*The Mixed Inference Problem.*¹¹⁷ For truth pluralism, the Mixed Inference Problem arises when we try to infer what seems to be a true conclusion from a set of premises that employ different notions of truth. Here is how we can make this worry more precise. According to standard accounts of semantic validity an argument is valid iff it necessarily preserves truth from the premises to the conclusion. A natural way of unpacking this idea is as follows: in a valid inference, there is some property – the property of *being true* – and an argument is valid iff this property is necessarily preserved from the premises to the conclusion.¹¹⁸ However, on a pluralist picture, where there exists more than one property that can make a proposition true, it becomes difficult to see what property is necessarily preserved throughout a valid inference. Here is how Tappolet states the worry.

“... there is a simple and equally powerful objection to the claim that there is a plurality of truth predicates. Consider the following inference:

- (1) Wet cats are funny.
 - (2) This cat is wet.
- Ergo, this cat is funny.

¹¹⁶ I am following Lynch (2009: 54) who presents Simple Alethic Pluralism - essentially a truth analog of (SJP) as an ambiguity view.

¹¹⁷ This objection to alethic pluralism was famously presented by Tappolet (1997).

¹¹⁸ Sometimes the idea is expressed in terms of predicates rather than properties – e.g. Tappolet (1997).

The validity of an inference requires that the truth of the premises necessitates the truth of the conclusion. But how can this inference be valid if we are to suppose with Crispin Wright that two different kinds of truth predicates are involved in these premisses? For the conclusion to hold, some unique truth predicate must apply to all three sentences. But what truth predicate is that? And if there is such a truth predicate, why isn't it the only one we need?" (1997: 209-210)

Now, it seems likely that we can expect an analogue problem to arise for (SJP).¹¹⁹ With two different notions of justification in play, we can expect (SJP) to run into problems with the validity of seemingly unproblematic inferences when the premises of an argument employ different notions of justification. Consider the following example.

You have invited Teddy and Bobby to a party. Teddy has already told you that if he comes to the Party so will Bobby because they *never* go to parties without the other person. Now, let's assume that Teddy is generally reliable about these sorts of things and that you have no reason to doubt the sincerity of his assertion such that your belief that 'if Teddy comes to the party, then so will Bobby' is *normally true*. So, you are justified_n in believing the conditional; $J_n(T \rightarrow B)$. Teddy then calls you and tells you that he is planning on coming to the party and that the only reason he would not attend is if he wins a very large lottery, which is drawn this afternoon. In that highly unlikely scenario he would go out for a celebratory dinner instead. While the belief that Teddy will come to the party is not true in all normal worlds, it is

¹¹⁹ In an influential paper, JC Beall (2000) provided a response to the Mixed Inference Problem on behalf of the truth pluralist. Beall argued that the standard account of semantic validity does *not* require that there is some unique *truth property* or *truth predicate* that is necessarily preserved from the premises to the conclusion of an argument. All that the traditional account of validity requires is necessary truth-preservation. And this more general account will be satisfied as long as there is *some* truth property or truth predicate that is necessarily preserved throughout an argument. Here is Beall, "The pluralist's reply, then is straightforward. Pluralists are committed to there being two different ways of being true. This however, does not conflict with the usual semantic account of validity. Validity is still necessary truth preservation; however, 'truth preservation' must be understood pluralistically – as the preservation of any way of being true." (2000: 382). This point was also made more recently by Wright (2013 "But to run this objection [Tappolet's and related versions], you first need to ground the idea that validity has to be preservation of a single property. And so far, as I can see, there is no independent motivation for that." (133). If we accept Beall's and Wright's argument, then neither alethic pluralism nor J-pluralism need to worry about the Mixed Inference Problem. This of course is good news for the prospects of J-pluralism.

nevertheless overwhelmingly probable and therefore justified_p; J_p(T). Since you are justified in believing that if Teddy comes to the party so will Bob, and you are justified in believing that Teddy is coming to the party, it seems that you should also be justified in believing that Bobby will come to the party. On (SJP) however, it is not clear how this inference is supposed to work, for what property would justify the belief that Q? The belief in the conditional is justified as a result of being normally true while the belief that Teddy will be at the party is justified as a result of being highly probable. But with two different notions of justification employed in the two premises - J_n(T → B) and J_p(T) - it is unclear how we could infer anything about the justificatory status of Q. After all regardless of what property we think should justify Q (presumably the property of being highly probable), there would be no single property that is preserved from the premises to the conclusion.

What the mixed inference problems shows is that as soon as we have more than one *kind* of justification, seemingly simple and unproblematic inferences fail. Importantly, since inferences are at the heart of rational deliberation, reasoning, and argumentation, the mixed inference problem presents a serious challenge for (SJP).

*The Mixed Compound Problem.*¹²⁰ Another problem for simple pluralism is the Mixed Compound Problem. For the truth pluralist the problem of mixed compounds is sometimes considered to be the most difficult challenge.¹²¹ Consider a conjunction of two propositions that are true in different ways:

- (1) Snow is white and killing babies is wrong.

Given that both conjuncts are true it seems convincing that the conjunction should also be true. But in what way? The first conjunct is true because it corresponds to the external world while the second conjunct is often taken to be true according to some notion of coherence. But, since the conjuncts are true in different ways it is not clear

¹²⁰ This problem was first introduced by Williamson (1994). See also, Tappolet (2000).

¹²¹ The reason the Mixed Compound Problem is often considered to be more difficult for the truth pluralist is that some moves that might get them out of the Mixed Inference Problem will not be available in the case of the Mixed Compound Problem. For our current purposes however, we do not need to consider this issue in more detail.

in what way the resulting conjunction would be true. Since the second conjunct does not correspond to anything in the external world, the proposition expressed by the conjunction cannot be true according to a correspondence notion of truth. And since the coherence notion that explains why the second conjunct is true will not apply to the first conjunct, which expresses a proposition about the external world, the conjunction cannot be true according to this coherence notion either.

One might of course worry that this problem for simple alethic pluralism will apply equally to simple pluralism about justification. More precisely, we might worry that if we accept (SJP), it will be difficult to explain what justifies the belief in a conjunction when the individual conjuncts are justified in different ways. Consider again the example of Teddy and Bobby.

We said that the belief that if Teddy comes to the party then so will Bob is justified because it is true in all normal worlds. Furthermore, we said that the belief that Teddy is coming to the party is justified because it is highly probable. Now, presumably we would also want to say that one is justified in believing the conjunction of these claims, namely that:

(2) If Teddy comes to the party, then so will Bobby, and Teddy is coming to the party

But *in what way* would the belief in this conjunction be justified? Put differently, what property justifies the belief in this conjunction? One might worry that since the two conjuncts are justified in different ways – i.e. by different properties - neither notion of justification will apply to the conjunction.

Fortunately, however, we defined epistemic environments in a way that allows us to avoid this problem. Recall that unlike the alethic pluralist we defined epistemic environments *exhaustively*, so that for any proposition we will either be in a normal epistemic environment or in a risk accumulating epistemic environment (but never both). Also recall that which epistemic environment we are in is determined by whether or not there exist any normal worlds in which we falsely believe that P. Hence, even for conjunctions like (2) we will be in a unique epistemic environment, which determines what property will make the belief justified.

In what epistemic environment are we with regards to (2)? To answer this question, we must ask whether there exists a normal world in which we falsely believe (2). Now, for the second conjunct there does exist a normal world in which we falsely believe that P; this of course is the world in which Teddy wins the lottery. Since a conjunction is false if any of its conjuncts are false and since there does exist a normal world in which we falsely believe one of its conjuncts there will also exist a normal world in which the conjunction is false. So, given our evidence we are in a risk accumulating epistemic environment with regards to (2). This means that (2) is justified as long as it is sufficiently probable on the evidence. So, it seems that pluralism about justification – even on the simple picture – can avoid The Mixed Compound Problem. Next, I want to consider one final objection to (SJP).¹²²

The Commonality Objection. When we say that a number of beliefs are justified, it seems somehow natural to construe what we are saying as something along the following lines, ‘there exist a number of beliefs that all share a common property, namely the property of *being justified*.’ Put differently, it seems natural to think that *being justified* is a somehow unified property – i.e. a property that all justified beliefs have in common. However, if we accept (SJP) according to which justification is ambiguous and picks out different properties (or concepts) in different contexts, the notion of justification will not be unified. We might take the inability to accommodate this idea of *commonality* amongst the set of justified beliefs to indicate that (SJP) is somehow on the wrong track.¹²³

¹²² For alethic pluralists The Mixed Compound Problem arises in much the same way for disjunctions. However, J-pluralist – even on the simple (SJP) interpretation - can run an analogous argument for dealing with disjunctions. Consider the following example.

(3) If Teddy comes to the party, then so will Bobby, and Teddy is coming to the party

What property would justify (3). Since a disjunction is only false if both of its disjuncts are false, we will need to determine whether there is a normal world in which both disjuncts are false. Since we said that the first disjunct is true in all normal worlds, there will be no normal world in which (3) is false. This means for (3), given our evidence, we are in a normal epistemic environment. Hence, what justifies the belief in (3) is the property of *being normally true*.

¹²³ Lynch (2009) points towards a related worry for simple versions of alethic pluralism, which will also hold for (SJP). The simple pluralist picture, according to which justification is simply ambiguous, more than anything is a pluralist picture of the *meaning of the term* ‘justification’. In other words, it is not actually a metaphysical pluralism but instead a purely semantic type of pluralism about the term ‘justification’. As Lynch (58) points out, even though the term ‘bank’ is ambiguous between two different meanings (the river bank and the financial institution) we do not therefore conclude that there

Thus, while the simple version of J-pluralism (SJP) is able to avoid the Problem of Mixed Compounds, there remain two serious challenges – The Mixed Inference Problem and the Commonality Objection. These worries should lead us to abandon (SJP). In the next section I propose a metaphysical framework for thinking about J-pluralism that can avoid these issues.

5. *The Functional Theory of Justification*

In the truth pluralism debate, problems like the Mixed Inference Problem and the Mixed Compound Problem have led to the general consensus that simple versions of alethic pluralism – truth analogues of (SJP) according to which the term ‘truth’ ambiguously picks out different properties (or concepts) in different contexts - are unsustainable.¹²⁴ Subsequently, truth pluralists have drawn the following conclusion about the structure of a viable metaphysics for alethic pluralism: truth must be both one and many.

The details of this proposal have been developed in different ways. For our current purposes I am going to focus on one very prominent one due to Michael Lynch.¹²⁵ Lynch (2009: 67-68) cashes out the idea that truth is both ‘one and many’ as follows.

Truth is One: there is a single property named by “truth” that all and only true propositions share.

is more than one way to be a bank or that we should be *bank pluralists*. In an important sense, (SJP) would be not a pluralist theory of justification at all since we would simply stop talking about justification and would instead start talking about justification_n and justification_p – following Lynch this would lead to a kind of metaphysical *nihilism* about justification rather than some form of metaphysical pluralism.

¹²⁴ Interestingly Lynch observes that no one seems to have actually defended this simple ambiguity version of alethic pluralism, but he notes that, “... a glance at the literature, together with anecdotal evidence, suggests that lots of folk seem to think alethic pluralists must be committed to it.” (2009: 54)

¹²⁵ While Lynch develops the idea that truth must be both one and many in terms of *one property* that can be realized by *many properties*, Wright (1992, 1999, 2013) has previously proposed and defended the idea that there is *one concept* that is satisfied by *many properties*. However, a detailed discussion of these issues would take us too far off course and is not immediately relevant to our current discussion.

Truth is Many: there is more than one way to be true.

How may a theory of truth accommodate these two constraints? To accommodate them Lynch (2009) introduces the Functional Theory of Truth.¹²⁶ According to Lynch's functional theory, *being true* is a functional property that can be realized in more than one way. On this picture, broadly speaking, there will be a single property – the property of *being true* - that is shared by all true beliefs. But, beliefs (or propositions) can come to acquire this property in more than one way.

Since, as we have seen, some of the problems for truth pluralism apply equally to J-pluralism, it seems reasonable that we should follow the truth pluralists' proposal and take on board the idea that justification must be both, one and many. Following Lynch, this yields the following two constraints, which a pluralist theory of justification needs to accommodate.

Justification is One: there is a single property named by “justification” that all and only justified beliefs share.

Justification is Many: there is more than one way for a belief to be justified.

If we are able to combine these two metaphysical principles, then it seems J-pluralism too will be able to avoid the Mixed Inference Problem and the Commonality Objection; there will be a single property – the property of *being justified* – shared by all and only justified beliefs, and at the same time we will be able to make good on the promise of pluralism as there will be more than one property in virtue of which a belief can come to be justified. In what follows I will develop a *Functional Theory of Justification* modelled after Lynch's Functional Theory of Truth.¹²⁷ The trick to pull

¹²⁶ See Wright (2013) for a good discussion of alternative ways of satisfying the two constraints.

¹²⁷ Again, I am following Lynch because currently his proposal is the perhaps most prominent one in the literature. It is important to point out however that J-pluralism will also be compatible with other pluralist metaphysics – e.g. Wrights 'one concept/multiple properties' picture or other versions of the 'one property/many properties picture (see Wright (2013)). So, anyone who does not like Lynch's

off is to explain how justification could at the same time be both one and many. First, I will provide an outline of Lynch's Functional Theory of Truth. I then apply the framework to J-pluralism.

5.1 Lynch's Functional Theory of Truth

Any time we set out to provide a functional theory of some target notion, we must first specify the target notion's function, or its role. People sometimes use the metaphor of providing a 'job description'. This is where the *Platitudes-First Approach* previously discussed in chapter 1 becomes relevant again.

The metaphysical backbone for the alethic pluralism defended by Wright and Lynch, as outlined in chapter 1, is a certain *minimalism about truth*. According to this minimalist picture, truth, contra deflationism, *is* a substantive property but its nominal essence is exhausted by a set of basic platitudes – i.e. a set of basic *a priori* laws or principles that chime well with how we ordinarily think about truth (Wright 1999: 226). Wright for instance identifies the following as truth's basic platitudes, which he argues exhaust the essential features of the truth predicate (1992: 34).

- that to assert is to present as true;
- that any truth-apt content has a significant negation which is likewise truth-apt;
- that to be true is to correspond to the facts;
- that a statement may be justified without being true, and vice versa;

For Wright, answering the question 'what is truth?' essentially reduces to the question of what predicate(s) can satisfy these platitudes. Lynch (2009, Ch. 1) follows a very similar strategy; he argues that in discovering the nature of truth, or truth's *nominal essence*, a plausible starting point is a certain set of truisms, or folk beliefs, that capture how we ordinarily think and talk about truth. Lynch (2009: 70) identifies the following as truisms about our ordinary notion of truth.

functionalism is free to choose their favourite pluralist metaphysics and to embed J-pluralism in that framework instead.

Objectivity: The belief that p is true if, and only if, with respect to the belief that p , things are as they are believed to be.

Norm of Belief: It is *prima facie* correct to believe that p if and only if the proposition that p is true.

End of Inquiry: Other things being equal, true beliefs are a worthy goal of inquiry.

In Lynch's Functional Theory of Truth it's these truisms, or core platitudes that provide truth's job description, or truth's functional role - the *truth-role*. For example, "portraying things as they are" is part of the 'truth role' in that it is something that propositions which have the property of truth do" (2009: 71). So, in order for a property to qualify as a truth predicate it must satisfy the truth-role. This allows Lynch to say the following about the truth concept (72).

(F) $(\forall x)$ x is true if, and only if, x has a property that plays the truth-role.

(F) leaves it open how many properties there are that play the truth-role. There could be just one such property that plays the truth-role, or there could be more than one. The alethic pluralist of course argues that there is more than one property that can play the truth-role. As Wright points out with respect to the platitudes, "there is no reason to expect that the minimal platitudes will constrain their interpretation to within uniqueness. A variety of predicates may qualify as truth predicates, and we have to be receptive to the possibility that the truth predicates across different regions of discourse may differ in important respect." (1992: 75) Similarly, since it is the platitudes that provide the truth-role, there is no reason to expect there that there will be only one property that satisfies the truth-role.

What does it mean for a property to play the truth-role? In order for a property to play the truth-role, i.e. to satisfy truth's core truisms, it needs to have certain features that allow it to play this role. In other words, in order to play the truth-role a property needs to have certain features; features that allow the property to satisfy the core truism Objectivity, Norm of Belief, End of Inquiry, etc. Lynch, calls these

features the *truish-features*. For instance, for any property T and proposition P, T plays the truth-role, i.e. is a truth predicate, if it has the following features: “P is T if, and only if, where P is believed, things are as they are believed to be; other things being equal, it is a worthy goal of inquiry to believe P, if P is T; it is correct to believe P if and only if P is T...” (2009: 72). In short, Lynch spells out the metaphor of ‘playing the truth role’ as having certain features which allow a property to play the truth-role.

Putting these ideas together provides the two main tenets of Lynch’s Functional Theory of Truth (73).

- T1. In order for a proposition to be true it must have a property that plays the truth role; and
- T2. A property plays the truth-role if it has the truish-features specified by truth’s core truisms.

How does this functional theory of truth meet the demand that truth must be one and many? First, the truth functionalist is able to say that there does exist a single property shared by all true beliefs or propositions. This is the property had by all beliefs, “just when things are as they are believed to be; had by beliefs at the end of inquiry and which makes propositions correct to believe.” (74) Just like *being coloured* is an essential feature of the property *being red*, so *having the truish-features* is an essential feature of the property *being true*. So, *being true* is the property, ‘that is, essentially, had by beliefs just when things are as they believed to be; had by beliefs at the end of inquiry and which makes propositions correct to believe.’ (74) Importantly, any and all true proposition will share this single property. In this sense truth is one.

Second, the Functional Theory of Truth is able to accommodate the demand that truth is many. This is the case because the functional theory of truth is compatible with there being more than one property that plays the truth-role. Thus, there may be different ways of realizing (or manifesting)¹²⁸ the truth property. In some domains

¹²⁸ What does it mean for a property to manifest another property? Recall that on Lynch’s picture, the truth realizing properties – e.g. correspondence or coherence – are ontologically distinct properties from the truth property itself. There is the property of *being true* and there are the properties that domain-specifically *manifest* truth. A property M manifests another property F, broadly speaking, iff all essential features of F are also features of M (Lynch 2009: 75). In this case everything that has the

truth will be realized by a proposition corresponding to the external world and in other domains truth will be realized by some notion of coherence. Thus, truth is both one and many. In short then, according to the Functional Theory of Truth, truth is a functional property and there is more than one way in which this property can be manifested. Next, I will apply this framework to construct a *Functional Theory of Justification*.

5.2 The Functional Theory of Justification

Following Lynch's strategy, we first need to specify justification's functional role – the *j-role*. The *j-role*, or justification's job description, will be given by justification's core platitudes introduced in chapter 1 – *Truth Candidacy*, *Permission*, and *Blamelessness*. For instance, 'being a good candidate for being a true belief' and 'being at least permissible to believe' is something that is true of beliefs which have the property of being justified. Thus, for a property to qualify as a J-property it need to play the *j-role*.

(J) $(\forall x)$ x is epistemically justified if, and only if, x has a property that plays the *j-role*.

Again, (J) leaves it open how many such properties exist. The key idea driving J-pluralism is of course that there exists more than one property that plays the *j-role*.

Now, in order for a property to play the *j-role* it must have certain features; features that allow it to play the *j-role*. In other words, to play the *j-role* a property must have features such that it only applies to beliefs that are good candidates for being true beliefs and that are permissible to believe, etc. Following Lynch's

property M will also, as a consequence of having M , have the property F . In this case, following Lynch's use of the term, M manifests F . In the case of truth then, since the truth-realizing properties (correspondence and coherence) have the truth-features specified by the truth specified by the core truths, every belief that has one of the truth-realizing properties will also have the property of being true. Thus, the truth-realizing properties will *manifest* the property of *being true*.

An important upshot of cashing the 'realization' relation out in terms of manifestation as Lynch suggests is that although truth will be a functional property, it will not become a higher-order property. So, truth despite being a functional property must not be a property of properties. Instead, it is a first-order property that is manifested by other properties.

convention, we may call them *Jish-features*. So, sticking to our three core truisms, for any property J and belief P, J plays the j-role, i.e. is a justification-property, if it has the following features: P has J only if, P is a good candidate for being a true belief; if P is J, then it is at least permissible to believe P; and if P is J, then if P turns out to be false one is blameless for believing P.

The following then are the core tenants of the *Functional Theory of Justification*.

- J1. In order for a belief to be justified it must have a property that plays the j-role;
and
- J2. A property plays the j-role if it has the Jish-features specified by justification's core truisms.

With this metaphysical picture in play, the functional theory of truth can meet the demand that justification is both one and many. In other words, this framework allows us to capture both (J-One) and (J-Many). How?

Justification is one, in the sense that there exists a single property - the property of *being justified* - that all justified beliefs have in common. More precisely, *being justified* is the property had by beliefs just when they are good candidates for being true beliefs, when they are permissible to belief, and when an agent would be blameless if the belief turned out to be false.

At the same time justification is many. It is many because the functional property of *being justified* can be realized (or manifested) in different ways. Sometimes justification is realized by the property of being true in all normal worlds and other times it is realized by being highly probable. What property justifies a belief in any given instance of course depends on the epistemic environment in which the belief was formed. If, for some proposition, we are in a normal epistemic environment – i.e. if our evidence is such that there exist no normal worlds in which we falsely believe that P - then justification is realized or manifested by the property of *being normally true*. On the other hand, if we are in a risk accumulating epistemic environment with regards to some proposition – i.e. if, given one's body of evidence, there does exist at least one normal world in which one falsely believes that P - then justification is realized by the property of *being highly probable*.

5.3 The Problems for Simple J-Pluralism Revisited

How does the *Functional Theory of Justification* solve the problems faced by simple version of J-pluralism (SJP)? Let's consider the two challenges in turn.

Mixed Inference Problem. Recall that simple versions of J-pluralism struggle to explain the validity of seemingly unproblematic inferences in cases where different premises of an argument employ different notions of justification. The root of the problem was that in these cases, there is no single J-property that is preserved from the premises to the conclusion – we called this the Mixed Inference Problem. The *Functional Theory of Justification* avoids this problem because there now does exist a single property that is shared by all justified beliefs. Even though the justificatory status of a belief might be realized by different properties, there nevertheless is a single property – the property had by beliefs just when they are good candidates for being true beliefs, when they are permissible to believe, and when an agent would be blameless if the belief turned out to be false - that is shared by all justified beliefs. This means that on the functionalist picture, even if justification is realized by different properties across different premises of an argument, there does exist a single property that is preserved from the premises to the conclusion of a valid inference.

The Commonality Objection. According to the commonality objection simple versions of J-pluralism, cannot accommodate the idea that there is a single unifying property that is shared by all justified beliefs. In other words, (SJP) cannot accommodate the plausible idea that there exists a single property that all justified beliefs have in common. The *Functional Theory of Justification* can avoid this problem. On the functional theory, there does exist a single property common to all justified beliefs. Again, this is the property had by beliefs just when they are good candidates for being true beliefs, when they are permissible to believe, and when an agent would be blameless if the belief turned out to be false. Since there does exist a single unifying property shared by all justified beliefs, the functional theory of justification can avoid the commonality objection.

Thus, the *Functional Theory of Justification* outlined above appears to provide a promising metaphysics in which to embed J-pluralism.

5.4 Localizing (MPC)

One thing left to do at this stage is to reformulate (MPC) so that it's consistent with J-pluralism and the functional theory. In its original formulation, (MPC) takes the form of a global principle, according to which any of one's justified beliefs are closed under multi premise deduction. However, on the functional theory proposed above, justification can be realized by different properties some of which are compatible with multi premise deduction (e.g. the property of being normally true) and some which are not (e.g. the property of being highly probable). This means that on the functional theory we will need to limit the scope of (MPC). On the functional theory, multi premise deduction will hold as *a general principle* for all beliefs formed in a Normal Epistemic Environment – i.e. multi premise deduction will hold as a general principle for all beliefs that are justified by the property of being normally true. In contrast, due to the problem of risk accumulation, multi premise deduction will not hold as a general principle for beliefs formed in a Risk Accumulating Epistemic Environment – i.e. multi premise deduction will not hold as a general principle for all beliefs that are justified by the property of being highly probable.¹²⁹ Here is a localized version of (MPC) that takes this feature of the functional theory into account.

Multi Premise Closure Local (MPCL) If the belief that p is justified in virtue of being true in all normal worlds, and the belief that q is justified in virtue of being true in all normal worlds, ... and the belief that n is justified in virtue of being true in all normal worlds, then the conjunction ($p \& q \dots \& n$) will be justified in virtue of being true in all normal worlds.

On the functional theory then we can preserve (MPCL) – a localized version of (MPC) - as a general principle.

6. Objection and Replies

¹²⁹ It is important to note that for beliefs formed in a Risk Accumulating Epistemic Environment multi premise deduction will still hold in *many* instances. It only fails as a general principle.

For the remainder of this paper I am going to anticipate objections to the pluralist theory of justification proposed above and offer some replies.

Objection 1. The perhaps most pressing worry for the pluralist account of justification presented above is that it might strike some readers as being ad hoc or too tailor-made for solving a particular problem - i.e. the *(MPC)-Stinginess Dilemma*.

Reply. I think that there are at least two possible responses to this objection. I will briefly consider them both.

One option would be to outright deny that there is anything ad hoc or tailor-made about J-pluralism. The metaphysical underpinning of J-pluralism is a larger, and independently motivated, minimalism about *justification*. As I argued in chapter 1, according to *The Minimalist Conception of Justification*, the nominal essence of justification is exhausted by a set of core platitudes. And any property that satisfies the core platitudes is a candidate justification-conferring property. From the minimalist conception there is of course a natural way into pluralism, for there may be more than one property that satisfies justification's core platitudes. With this methodological picture in place, there is nothing *ad hoc* or even surprising about a pluralist theory of justification. So, the pluralism I defend should be seen as standing on a solid and independently motivated methodological foundation.

At this point an objector might grant that the *ad hocness* of J-pluralism does not enter at this early stage, but instead enters later when we attempt to spell out the view in more detail. For instance, someone might grant that the motivations underwriting J-pluralism are fine, but that the way in which we defined the different epistemic environments is ad hoc. But again, I think there are good reasons to resist this accusation. The fact that sometimes our evidence is such that there is no normal world in which we falsely believe that P while at other times our evidence is such that there do exist normal worlds in which we falsely believe that P is a compelling and arguably important distinction; we can even state the distinction in precise modal terms. What, one might ask, is ad hoc about it? Recall, the normalcy views of justification defended by Smith (2010, 2016) and Leplin (2009) rest precisely on this

distinction. It is difficult to see why there should be anything ad hoc about taking this distinction seriously.

So, if the objection that the pluralism defended here is ad hoc is supposed to stick, then someone would have to show precisely *what* feature of the account is ad hoc and *why*. Moreover, as we have seen, the pluralist picture of justification I defend closely mirrors the kind of pluralism previously proposed by alethic pluralists. As a result, we might think that any argument intended to show why pluralism about justification is ad hoc will easily translate over to truth pluralism. But if this is the case, then the objection would just be a general objection to pluralist projects in philosophy rather than one that specifically targets the justification pluralism I defend. In this case, the objection, if it indeed sticks, should not be counted as a direct strike against J-pluralism but against pluralism in general. In any case, more work would need to be done to provide a compelling reason for thinking that the pluralist picture of justification outlined above is ad hoc or in some way unprincipled.

A second option for responding to the ad hocness accusation (one that I find less attractive) would be to concede that maybe there is something ad hoc or tailor-made about the pluralist theory of justification proposed here but deny that this makes for a forceful objection. We started out with a dilemma for monist theories of justification: if we opt for a probabilistic theory of justification, then we need to give up (MPC) and if we opt for a non-probabilistic theory which can preserve (MPC), then justification becomes too stingy. Since both the preservation of (MPC) and not being too stingy are plausible desiderata for a theory of justification, monism is in trouble. Now, on a pluralist theory we can accommodate both desiderata – we can preserve a version of (MPC) - viz. (MPCL) - without justification becoming too stingy. As a result, pluralism about justification is able to capture more of the desiderata we would like a satisfactory theory to capture than its monist counterparts. Now, even if we accept that this solution is in some sense tailor-made (whatever this means) to deliver this result, then there remains a question about what exactly the force of this objection is supposed to be. As long as J-pluralism captures more desiderata than its monist counterparts it seems that it is a preferable and worthwhile theory.

Objection 2. According to a second, closely related, objection one might think that opting for a considerably more complex pluralist theory of justification, *only* to preserve a version of (MPC) and to avoid some stinginess concerns is *too drastic*. In short: one might worry that in this case the cure is worse than the disease.

While some might find this an initially compelling objection, there are again a number of ways of responding to it.

First, we could push back against the attempt to downplay the considerable theoretical pay-off offered by J-pluralism. Ever since the 1960s, when the viability of (MPC) became a prominent concern in the justification literature, people have gone through considerable troubles to try and defend the principle. However, ultimately none of them have been particularly successful. It was only when philosophers like Smith and Leplin started exploring non-probabilistic approaches to justification that (MPC) once again appeared to be within reach. But, as we have seen, one consequence of preserving (MPC) is that justification will become very stingy. Against this background, the fact that pluralism about justification is compatible with deductive closure for a great many beliefs whilst at the same time being able to avoid any stinginess concerns should be seen as a *considerable achievement*. It allows us to preserve a principle that philosophers have unsuccessfully tried to defend for decades without incurring any of the significant costs commonly associated with preserving this principle. This is by no means a negligible achievement!

An alternative response would be to question what precisely the force of the alleged objection is supposed to be – i.e. why should we think that the cure is worse than the disease? One way of precisifying the objection might be by appeal to simplicity as a theoretical virtue. Since pluralism is arguably a more complex theory (I am not actually convinced that this is the case), so the objection might go, we should be reluctant to accept it unless we have *good reasons* for preferring it over its simpler competitors. In a next step, the objector might argue that preserving a version of (MPC) and avoiding stinginess worries are not sufficient pay-offs to make the theory preferable to its monist counterparts. Let's assume for our current purposes that this is the best way of spelling out the objection.

Again, there are a number of things we can say in response. First, we could of course reiterate the substantive achievement presented by a theory of justification that can accommodate (MPC) without being too stingy. Since monist theories cannot accommodate both desiderata, we may reject the second step of the argument, which attempts to downplay the considerable theoretical benefits offered by the pluralist theory of justification.

However, there is a second and perhaps more compelling response available to us: we could simply reject that the appeal to simplicity provides a compelling objection against justification pluralism. On what grounds might we deny this?

First, the appeal of simplicity is usually reserved as criteria for theory-choice in cases where two theories are *equally successful in their predictions*; in cases where two theories are predictively coextensive, so the thought goes, we might have a reason for preferring the simpler theory. In our case however monist theories of justification and the pluralist theory I defend are not equally successful in their predictions - the pluralist theory is more successful and capturing our desiderata than any of its monist competitors. So, the conditions under which simplicity becomes an appropriate criterion for theory-choice don't seem to be met in this case.

Furthermore, when we speak of 'simplicity' as a criterion for theory choice, we usually mean *ontological simplicity*. What does this mean? Broadly speaking, a theory A is simpler than theory B if A requires fewer *ontological commitments* than B. However, the difference between the pluralism I defend, and monist accounts of justification are not ontological. Pluralist accounts of justification do not require the existence of any additional entities. So, it is not clear that the simplicity argument will gain any traction against *The Functional Theory of Justification*.

A final response to the worry that pluralism about justification is perhaps too 'drastic' of a move, is to point out that pluralism about justification is not unprecedented. In the past a number of people have proposed pluralism about justification as a strategy for overcoming the apparent stalemate between internalists and externalists. With neither side being able to make a decisive case against their opponent, many epistemologists, particularly in the second half of the 1980s, suggested that there may simply be more than one concept of justification – one internal and one external (Alston 1985, Lehrer 1986, Goldman 1988). Goldman

(1988) for example introduces the two concepts of *strong* and *weak justification*. *Strong justification*, according to Goldman, applies to beliefs that are formed through a reliable belief forming process, whilst *weak justification* picks out some evidence-centered internalist notion of justification. Goldman of course assigns priority to beliefs that are formed using reliable methods, i.e. strongly justified beliefs, and characterizes weakly justified beliefs as “ill-formed-but-blameless” (Goldman 1988: 56). However, having two concepts of justification in play enabled Goldman to respond to many of the notorious counterexamples to his process reliabilism, like the New Evil Demon Problem: Goldman could now say that an agent in a new evil demon world, whilst lacking *strong justification* for their beliefs, can nevertheless have *weak justification*.

As mentioned above, other epistemologists have, for similar reasons, drawn similar conclusions. Alston (1985) distinguished between *deontological justification* and *evaluative justification*, where deontological justification is an internal notion tied to normative concepts like *obligations* and *blame*, while evaluative justification is an external notion connected purely to the goal of believing truths and avoiding falsehoods. Similarly, Lehrer (1986) distinguished between *personal justification* and *verific justification*. Personal justification, according to Lehrer, is an internal concept of justification while verific justification is an external concept of justification. The primary motivation for these pluralist views of justification, to reiterate, was the prospect of overcoming the apparent stalemate between epistemic externalists and internalists.¹³⁰

It is important to highlight that these proposals were never worked into full theories. No one tried to provide an overarching metaphysics outlining how the different notion of justification are meant to relate and interact with one another. And in their original formulations these proposals are best thought of as instances of Simple J-Pluralism (SJP), which we argued earlier, faces several problems. Nevertheless, these multi concept responses to the externalism/internalism debate are relevant for the pluralism I defend in at least two ways.

¹³⁰ For an overview of these pluralist approaches to justification, see Pedersen (2013, sect. 2 & 3).

First, it sets an important precedence for pluralism about justification in cases of a stalemate between two opposing, yet well-motivated, theories. Assuming that both the Lockean View and the Normalcy View are well-motivated theories of justification, then we shouldn't think that the pluralism I defend is anymore 'drastic' than the pluralism previously defended by Goldman, Alston, and Lehrer.

Second, the multi concept response to the externalism/internalism debate provides an illustrative example of how two accounts of justification may be combined to yield a stronger overall theory. Setting aside issues over the names *weak* and *strong* justification, Goldman's move towards a pluralist picture of justification allowed him to provide an overall more compelling theory of justification. Now note that the functional theory I propose is similarly motivated: equipped with two justification-conferring properties that apply in different epistemic environments, we can sidestep the *(MPC)-Stinginess Dilemma* and thereby overcome some of the major drawbacks of Lockean views and normalcy views. As we can see, the idea that there may be more than one way for a belief to be justified does have some relevant precedence in the literature; and, seeing that people have recognized the benefits of pluralism about justification - or J-pluralism - in the past, there is no reason to dismiss it too easily in this case.

Objection 3. The pluralist account is subject to over-determination worries. Many propositions that are justified as a result of *being normally true* will also be *highly probable*. In these cases, it is not clear what the justification-conferring property is.

Reply. First, it is true that members of the two families of views - the Lockean View and the Normalcy View - will deliver the same prediction in many cases. This however should not surprise us. If two theories of justification made vastly different predictions about a large number of cases, then this would suggest that one of the theories is on the wrong track and not a compelling theory of justification in the first place. So, it should not come as a surprise that many beliefs which are normally true will also be highly probable and *vice versa*.

Importantly however, the two theories are not coextensive. There will be beliefs that are highly probable for subject but fail to be true in all normal worlds (e.g.

lottery beliefs, and beliefs in cases like Allergy, Lightning, and Pill). Likewise, there will be beliefs that are normally true on the evidence but that are highly improbable (e.g. long conjunctions).

Despite the fact that many beliefs that are normically supported by the evidence will also be highly probable on the evidence it is not the case that this will lead to a problem of over-determination. For any given belief, the property that justifies the belief is determined by the agent's epistemic environment. An agent's epistemic environment in turn is determined by their body of evidence. If there is no normal world in which, on the evidence, S falsely believes that P, then S is in what I called a *Normal Epistemic Environment* with regards to P. And in a Normal Epistemic Environment justification is realized by the property of being true in all normal worlds. Alternatively, if there does exist at least one normal world in which S, given the evidence, falsely believes that P, then the agent is in what I called a *Risk Accumulating Epistemic Environment*. And in a Risk Accumulating Epistemic Environment a belief is justified iff it is sufficiently probable on the evidence. Since for any proposition and any body of evidence the subject can only be in one of the two epistemic environments – either there does exist at least one normal world in which the subject, given the evidence, falsely believes that P or there doesn't - it is impossible to be in two different epistemic environments with regards to the same proposition. It follows that that each proposition can be justified *in one and only one way*. Hence, the Functional Theory of Justification I defended can avoid the over-determination objection.

Chapter 6

Applications: New Solutions to Our Epistemic Puzzles

1. *Putting the Functional Theory to Work*

In the previous chapter I introduced a new pluralist theory of epistemic justification - the *Functional Theory of Justification* - and argued that it can help us solve a difficult problem faced by monist theories of justification: the (MPC)-Stinginess Dilemma. In this final chapter I argue that aside from providing a way out of the dilemma, the functional theory also provides interesting new solutions to a number of epistemic puzzles and paradoxes including the Lottery Paradox, the Preface Paradox, the Paradox of the Pill, as well as the previously considered puzzles involving statistical evidence. As we will see, the functional theory provides new resources for solving these puzzles in a way that avoids many of the costs incurred by its monist competitors.

The upshot of this is twofold. First, it provides further reasons in favour of the functional theory of justification. And second, it provides additional *independent* motivation for justification pluralism – i.e. it provides motivation that is entirely independent from the *(MPC)-Stinginess Dilemma*, which in the previous chapter was presented as the primary motivation for going pluralist.

2. *The Lottery Paradox*

As we have already seen in chapter 3, the lottery paradox is usually considered to put pressure on the claim that the Lockean picture of justification - which explains justification in terms of high probability - is compatible with the intuitively compelling and generally attractive principle of multi premise closure (MPC). The

lottery paradox does this by showing that the Lockean View in conjunction with (MPC) violates the plausible No Contradiction principle (NC), which states that one cannot justifiably believe a contradiction. We can state the challenge posed by the lottery paradox as follows: as long as one takes (NC) to present a fundamental constraint on epistemic justification, any theory of justification on which one can justifiably believe of any ticket in a sufficiently large lottery that it is going to lose – let’s call this a *lottery belief* - will be incompatible with (MPC).

Let’s briefly reiterate why this is the case using again as an example Juliet, who just bought a ticket for a very large, fair, lottery with a guaranteed winner. In chapter 3, we presented the lottery paradox as follows.

- (L1) For any ticket Juliet is justified in believing that the ticket will lose; $JB_J(t_1) \& JB_J(t_2) \dots \& JB_J(t_n)$.
- (L2) If Juliet is justified in believing of any given ticket that it will lose, then Juliet is justified in believing that all tickets will lose; $JB_J(t_1 \& t_2 \dots \& t_n)$.
- (L3) Since Juliet knows that the lottery is fair and that it has a guaranteed winner, Juliet is justified in believing that at least one ticket will win; $JB_J \sim (t_1 \& t_2 \dots \& t_n)$.

However, another application of (MPC) to (L2) and (L3) yields the following violation of (NC).

- (L4) Juliet is justified in believing that every ticket will lose and that one ticket will win; $JB_J((t_1 \& t_2 \dots \& t_n) \& \sim (t_1 \& t_2 \dots \& t_n))$.

As we have seen in chapter 3, the perceived view is that solving the lottery paradox will involve either a rejection of (L2) and thereby (MPC), or, alternatively a rejection of (L1) the idea that lottery beliefs can be epistemically justified.

Proponents of the Lockean View will need to accept the first option and deny (MPC)¹³¹, whilst proponents of non-probabilistic accounts of justification like the Normalcy View will opt for the second option and deny that lottery beliefs are epistemically justified. As Smith and Leplin have compellingly argued, for lottery beliefs there always exists at least one normal world in which one falsely believes that one's lottery ticket is a loser – this of course is the world in which one's ticket happens to be drawn. So, on a traditional monist picture of justification, solving the lottery paradox requires that we choose between (MPC) and justification in lottery beliefs; but we cannot have both. Since many find (MPC) as well as the idea that lottery beliefs can be justified attractive, we have reason to be unhappy with the solutions available to us on monist theories of justification. As I argue next, pluralism about justification offers a more attractive solution to the lottery paradox.

How would the Functional Theory deal with the lottery paradox? First, let's determine what epistemic environment we are in with regards to lottery beliefs. Since for any ticket in a large lottery there exists a normal world in which, given the evidence, we falsely believe that the ticket is a loser, we are in a Risk Accumulating Epistemic Environment with regards to lottery beliefs. And in a Risk Accumulating Epistemic Environment the property that makes a belief justified is the property of being highly probable. Thus, since in a sufficiently large lottery the belief that any ticket is going to lose is highly probable, we will be justified in believing of any given ticket that it is going to lose. As a result, on the pluralism I defend we can accept (L1) and with it the intuition that we can justifiably believe lottery beliefs. Does this mean we need to deny (MPC)?

On the pluralist picture I defend the answer to this question is slightly more nuanced than it is within a monist framework. On the functional theory, accepting (L1) does not force us to reject multi premise deduction outright. Instead we only need to reject multi premise deduction for *one* of the ways in which a belief can come to be justified, namely for justified beliefs formed in a Risk Accumulating Epistemic Environment where justification is manifested by the property of *being highly*

¹³¹ In chapter 3 we considered proposals to save (MPC) by modifying the Lockean View so that it rules out justification in lottery beliefs – we called this the *modificationist strategy*. However, as we have seen this strategy is ultimately unsuccessful.

probable. But, this does not mean that closure *in general* fails. For *any* belief formed in a Normal Epistemic Environment – e.g. ordinary beliefs based on perception, testimony, memory, etc. - in which a belief is justified when it is *normally true*, justification will continue to be closed under multi premise deduction.¹³²

As a result, on the pluralist picture I defend we can retain the justificatory status of lottery beliefs (L1) without having to outright deny that justification is closed under multi premise deduction. All that we need to accept is that for *one way* in which a belief can come to be justified closure fails. But this should not surprise us. Hence, the functionalist theory seems to provide a more attractive solution to the lottery paradox than its monist counterparts.

3. *The Paradox of the Pill*

The Paradox of the Pill introduced in chapter 3 – essentially a hybrid of the lottery and the preface paradox – gives rise to many of the same problems as the lottery paradox. And, as we will see, for many of the same reasons pluralism will again offer a more promising solution to the paradox. Consider again the following case

Pill S is given a bitter pill that ensures that a very small portion of S’s ordinarily justified beliefs, let’s say 1 out of every 10,000, chosen at random, will be false. The pill achieves this result by occasionally impairing S’s cognitive connection to the evidence resulting in occurrences of, for instance, misperceptions or false memories. Importantly however these occurrences are incredibly rare. Finally, S knows about the effects of the bitter pill.

From Pill we derived the Paradox of the Pill.

¹³² Recall that in the last chapter we argued that on the functionalist theory we can retain the following version of (MPC).

Multi Premise Closure Local (MPCL) If the belief that P is justified in virtue of being true in all normal worlds, and the belief that Q is justified in virtue of being true in all normal worlds, ... and the belief that N is justified in virtue of being true in all normal worlds, then the conjunction (P & Q ... & N) will be justified in virtue of being true in all normal worlds.

- (BP1) S, after taking the pill, continues to be justified in believing the propositions made sufficiently probable by her evidence; $JB_S(p_1) \& JB_S(p_2) \dots \& JB_S(p_n)$.
- (BP2) If S continues to form justified beliefs, then S is justified in believing the conjunction of her individually justified beliefs; $JB_S(p_1 \& p_2 \dots \& p_n)$.
- (BP3) S is justified in believing that the conjunction of her justified beliefs will contain at least some false belief; $JB_S \sim (p_1 \& p_2 \dots \& p_n)$.

However, taken together (BP2) and (BP3) violate the No Contradiction Principle (NC), as they give rise to the following contradiction.

- (BP4) S is justified in believing a contradiction, namely that the conjunction of her justified beliefs is true and that some members of the conjunction are false; $JB_S((p_1 \& p_2 \dots \& p_n) \& \sim (p_1 \& p_2 \dots \& p_n))$.

The options we have for responding to this paradox are essentially the same as in the lottery paradox. Either we reject (BP2) and thereby (MPC) or we deny (BP1) and insist that after taking the pill none of S's beliefs are justified.

Once again, proponents of the Lockean View will likely need to opt for the first option and deny (BP3). And, as was the case in the lottery paradox, non-probabilistic views like the Normalcy View will deny (BP1). Why? As we argued in chapter 4, after knowingly taking the pill, it seems reasonable to think that S's total body of evidence will be such that for any belief there exists a normal world in which the belief is false – this is the world in which the belief happens to be affected by the pill. Hence, none of S's post-pill beliefs will be justified. But, as was argued previously, this conclusion seems unpalatable. Responding to the fact that 0.001% of one's beliefs are going to be false by denying justification for all of one's beliefs (bracketing beliefs in necessary truths) and by embracing global skepticism about justification seems like an excessively costly and unconvincing response to the paradox. Hence, by forcing us to choose between two unattractive options, monism once again leaves us in between a rock and hard place.

Let's see how the functional theory I propose will deal with The Paradox of the Pill. First, we need to determine in what epistemic environment the agent is with regards to their post-pill beliefs. Recall that after taking the pill, for any belief that P, there exists a normal world in which S falsely believes that P. This means that after taking the pill, S will be in a Risk Accumulating Epistemic Environment for all beliefs. Since in a Risk Accumulating Epistemic Environment justification is realized by the property of *being highly probable*, S's post-pill beliefs will continue to be justified as long as they remain highly probable. And since for each belief there exists only a 0.001% chance that it was affected by the pill, we can expect that the majority of beliefs that would have been justified prior to S taking the pill will continue to be justified even after taking the pill.¹³³ Thus, the functional theory delivers the intuitively correct result concerning the justificatory status of S's individual beliefs. Seeing that in a Risk Accumulating Epistemic Environment, in which beliefs are justified if they are highly probable, justification is not closed under multi premise deduction, the functional theory will resolve the Paradox of the Pill by rejecting (BP3).

However, on the functional theory, unlike on the Lockean View, we don't need to abandon (MPC) wholesale. As argued previously, we only need to deny that for *some* beliefs – i.e. beliefs formed in a Risk Accumulating Epistemic Environment – justification can *sometimes* fail to be closed under deduction. One of the key advantages of the functional theory is that we can keep deductive closure as a general principle for all beliefs formed in a Normal Epistemic Environment.¹³⁴ In other words, on a pluralist picture solving the paradox only requires the concession that for *one* of the properties that can justify beliefs (e.g. the one that applies in risk accumulating epistemic environments) multi premise deduction fails. This concession however is perfectly compatible with the idea that for a different justification-conferring property

¹³³ There may of course be some exceptions. Consider a belief with a degree of probability just above the threshold required for justification. Now it is possible that for such a *borderline justified belief* the additional risk introduced by the pill – despite the fact that it is miniscule – will be sufficient to bring the degree of probability just below the threshold required for justification. Hence, even on the Lockean View and the Functional Theory the pill may affect the justificatory status of some beliefs. But I assume that these cases will be rather rare and that even if there are some such cases the result would still be nowhere near as severe as on the Normalcy View, where one loses justification for all beliefs.

¹³⁴ As we argued in the previous chapter we can keep (MPCL) as a general principle. See the previous footnote for the formulation of the principle.

(e.g. the property of *being normally true*) multi premise deduction will hold. Hence, the functional theory again seems to offer a more attractive solution to the paradox than its monist competitors – we can accommodate the justificatory status of the individual conjuncts and also keep a version of (MPC) - viz. (MPCL).

4. *The Preface Paradox*

Let's turn our attention to the preface paradox. Here is how we presented the paradox in chapter 3.

An author has just finished an ambitious, yet carefully researched, book. On the basis of her research the author has good evidence for every claim made in the book and is justified in believing that each claim made in the book is true. However, the author is also aware that when it comes to ambitious books, even the best scholarship can turn out to contain errors and in the past this has always been the case. In light of this the author includes in the preface of her work an apology for any errors in her book.

We captured the formal structure of the paradox as follows.

- (P1) The author is justified in believing each claim made in the book; $JB_A(p_1) \& JB_A(p_2) \dots \& JB_A(p_n)$.
- (P2) If the author is justified in believing each claim made in her book is true, then the author is justified in believing that all claims made in her book are true; $JB_A(p_1 \& p_2 \dots \& p_n)$.
- (P3) The author is justified in believing that her book contains at least one error; $JB_A \sim (p_1 \& p_2 \dots \& p_n)$.

However, taken together (P2) and (P3) combine to produce the following contradiction.

(P4) The author is justified in believing a contradiction, namely that all claims in her book are true and that her book contains at least one error; $JB_J((p_1 \ \& \ p_2 \ .. \ \& \ p_n) \ \& \ \sim (p_1 \ \& \ p_2 \ .. \ \& \ p_n))$.

How is the paradox usually resolved? A satisfactory response to the paradox is usually taken to require a rejection of either (P2), the claim that the author is justified in believing that her book is error-free, or (P3) the claim that the author is justified in believing that the book contains at least one error.

Since (P2) is the result of applying (MPC) to the individual claims in the book, rejecting (P2) would amount to a denial of (MPC). This is the solution that proponents of probabilistic accounts of justification following the Lockean View will need to endorse.¹³⁵ In contrast, proponents of non-probabilistic accounts of justification like the Normalcy View - which are committed to (MPC) - will deny (P3), the claim that the author is justified in believing that her book contains any errors. Why? On the

¹³⁵ In chapter 3 we considered some popular proposals by proponents of (MPC) for solving the paradox by denying (P3) instead. Central to these proposals was the notion of a *special reason*. For instance, Kaplan (2013) argued that the author is justified in believing that her book is error-free as long as there is no *special reason* for thinking that it contains any errors. A similar remark can be found in Leplin (2009), 'Notice how strange the inductive reason is. It is supposed to justify believing the negation of a conjunction without providing a reason or evidence against any conjunct... And these reasons are supposed to prevail over any justifications, however strong, however individually unimpeachable, that [the individual conjuncts] enjoy. I don't think so.' (98) However, there are a number of reasons to be suspicious of this strategy.

First, Kaplan leaves open what exactly a special reason amounts to and making this notion precise can reasonable be expected to be difficult. Second, we can strengthen the preface narrative to make the idea that the author lacks justification for believing that their book is error free increasingly implausible. This is precisely what Christensen (2004) does; he introduces additional reasons for thinking that the book contains an error – e.g. the author has published several books in the past all of which were found to contain at least one error, or that in the past even more meticulous scholars have been unable to write an error-free book, etc. Now, proponents of the special reasons defense can, and in fact do, continue to insist that none of this provides a special reason for believing that this book by the author contains any errors; Kaplan (2013) suggest the following, “[the author] should stand by her book – she should not shrink from saying that everything in it is true until she has found special reason to think it contains an error... We find it perfectly appropriate that an author should stand behind her book even in the face of the high probability that it, on account of it being a book of significant ambition, contains some error – even if there are other authors of comparably ambitious works, no less careful) and perhaps even more meticulous) than she, who have nonetheless failed to write error free books. (16) But, as we continue to strengthen the case against the claim that the book is error-free this strategy is beginning to look increasingly unpersuasive.

Most importantly however, the strategy fails because, as the Paradox of the Pill introduced in chapter 3 shows, we can setup puzzles that put pressure on (MPC) in which we do have a *special reason* for believing that the conjunction is false. In other words, the case against (MPC) does not require the absence of a special reason. As such, the *special reasons* strategy is not sufficiently general to get (MPC) out of trouble.

Normalcy View in order for a belief that P to be justified, there cannot exist a single normal world in which S falsely believes that P. It follows that if the author is justified in believing each of the claims in the book, then there will also be no normal world in which the author falsely believes the conjunction of the individual claims; for this would require a normal world in which one of the conjuncts are false and this is ruled out by normalcy views. Thus, on normalcy views the author will be justified in believing that her book is error-free.¹³⁶

What about the author's belief that the book contains at least one error? Following Smith (2016), I think that it is reasonable that for this belief there does exist at least one normal world in which the belief is false - this of course is the world in which the author's careful research has paid off and all the individually justified beliefs in the book are in fact true. Hence, despite the fact that the author's belief that her book contains at least one error is overwhelmingly probable, normalcy views are likely to predict that this belief is unjustified.¹³⁷

The solution offered by the Normalcy View is of course not without problems. Proponents of the Lockean View, or, more generally, anyone who thinks that the correct solution to the preface paradox is to reject (P2) – and as a result (MPC) - is likely to find the result that the author is justified in believing that her book is error free unconvincing. For instance, according to Christensen's (2004) preferred analysis, the preface paradox presents us with a choice between *epistemic modesty* on one hand

¹³⁶ To illustrate this point, it may be helpful to briefly consider Smith's precisification of normality via the notion of *calling for an explanation* (we could equally use Lepin's way of spelling our normality via the notion of *natural presupposition*). Consider whether it would call for a special explanation if it turned out that the author's belief that her book is error-free turned out to be false. It seems reasonable to think that it would. After all, the author carefully researched every claim made in the book and only included those proposition which were well supported by her evidence. If her book nevertheless turns out to contain errors, then this surely would require some explanation - mitigating or interfering circumstances would have to be at play to explain why, despite meticulous research, one or more of the claims turn out to be false.

¹³⁷ We can again illustrate this point using Smith's notion of *calling for an explanation*. Just consider whether it would call for any special explanation if the author's belief that the book contained at least one error was false. It seems reasonable to think that it wouldn't. After all, the author carefully researched every claim made in the book and only included those proposition which were well supported by her evidence. Thus, given the evidence, we should almost expect that every claim in the book is true, and subsequently, that the claim that there is an error in the book is false. In any case, given the evidence, it surely wouldn't be in any sense metaphysically abnormal if the belief that belief that the book contains errors was false.

and *epistemic immodesty* on the other. He calls the belief that the book contains at least one error the *Modest Preface Proposal* (MPP) and the belief that the book is error-free the *Immodest Preface Proposal* (IPP). He then goes on to argue that the kind of immodesty promoted by those who hold that author is justified in believing that her book is error free despite the strong inductive evidence against this claim, is unreasonable. Instead, Christensen argues we should embrace *epistemic modesty* and accept that the author is justified in believing that her book contains at least one error. What should we make of this way of analysing what's going on in the preface paradox?

Christensen's argument of course pivots on the idea that it is somehow unreasonable to believe something that is highly improbable. As Kaplan points out, Christensen argument is implicitly appealing to the following principle, 'you should not be willing to say that P in the context of inquiry if you are aware that P is improbable.' (Kaplan 2013: 19) But of course this principle is precisely what proponents of the Normalcy View (an other's sympathetic to (MPC)) will deny. So, while we may appreciate the force of Christensen's argument it seems that analysing the preface paradox in terms of modesty versus immodesty does little more than trying to tilt the debate in Christensen's favour by using normatively loaded language – after all who would want to admit to being epistemically immodest? However, anyone who wants to respond to the paradox by preserving (MPC) – like proponents of the Normalcy View - will simply reject that the author is being immodest or that the alleged immodesty is in any way problematic. Instead, they are likely to argue that there is nothing immodest about an author standing by their book in the absence of any *special reason* for thinking that one of its claims is false. It appears then that we have reached a stalemate between proponents of (MPC), who want to solve the paradox by rejecting (P3), and those who favour epistemic modesty and think that we should solve the paradox by rejecting (P2).

What does the functional theory of justification have to say about the preface paradox? Let's continue to assume that for every claim made in the book, there exists no normal world in which that claim is false. This seems like a reasonable assumption seeing that the author has carefully researched every claim in the book and been meticulous about avoiding any errors. If this is the case, then for every claim, the

author is in a Normal Epistemic Environment and each claim in the book will be justified in virtue of it being normally true. Since there does not exist a normal world in which any of the claims in the book are false, it follows that there will also be no normal world in which the conjunction of these claims – i.e. the claim that the book is error-free – is false. As a result, for the claim that the book is error-free the author will also be in a Normal Epistemic Environment; and since the claim is true in all normal worlds, the author has justification for believing it.

Now, what about the author's claim in the preface that the book contains at least one error? Based on what we have said above, it appears that there does exist a normal world in which this claim is false. This is the world in which the meticulous research of the author paid off and all the claims in the book are true.¹³⁸ Hence, with regards to the proposition that book contains at least one error the author is in a Risk Accumulating Epistemic Environment where a belief that P is justified as long as it is highly probable. And since the proposition that the book contains at least one error is highly probable – after all the inductive evidence strongly supports this claim - the author is justified in believing it.

In light of these considerations, the functional theory of justification offers what appears to be a novel and interesting *diagnosis* of what is going on in the preface paradox. In the preface paradox the two ways in which a belief can be justified are in conflict: the belief that the author's book is error free appears to be justified in virtue of *being normally true*, while the belief in the negation, namely that the book contains at least one error, is justified in virtue of *being highly probable*. In other words, in the preface paradox the two different ways in which a belief can be justified are pulling us in different directions. We have two incompatible beliefs that both seem to be justified (in different ways). This diagnosis provides a ready explanation for why we feel so strongly inclined to accept both (P2) and (P3) of the preface paradox: both beliefs have justification-conferring properties (or J-properties).

So much for a diagnosis; now how can proponents of the functional theory respond to the paradox?

¹³⁸ For a more detailed explanation of this see, the previous footnote.

4.1 Reduction to a More General Question

On the functional theory of justification, the preface paradox is not simply a puzzle about whether we should be epistemically modest or immodest. Instead, on the functionalist theory solving the preface paradox reduces to the following more general question: how should we deal with cases in which the two different ways of being justified lead to contradictions? In other words, responding to the preface paradox reduces to the more fundamental question of how to deal with cases in which the belief that P is true in all normal worlds, but the belief that $\sim P$, while failing to be true in all normal worlds, is nevertheless highly probable (and *vice versa*). As such, answering this question will ultimately require us to answer the undeniably difficult question of how to weigh the two different justification-conferring properties against each other. This will undeniably be a difficult task and it is unlikely that I will be able to provide a definite answer here. Fortunately, however, the functional theory by itself does not require us to take an official stance on this issue – as we will see, the functional theory is compatible with a number of different proposals. In what follows I discuss three options for dealing with conflicts between the two ways of being justified and outline the solution to the preface paradox that falls out of each of them. Finally, despite the fact that I will remain officially neutral between these options, I am going to provide some reason in favour of one of them.

4.2 Three Options

First, we could assign equal weight to beliefs formed in a Normal Epistemic Environments – i.e. to beliefs that are true in all normal worlds - and to highly probable beliefs formed in a Risk Accumulating Epistemic Environment. According to this proposal, the different ways of being justified would be *epistemically on par*.

Parity Beliefs that are true in all normal worlds and beliefs that fail to be true in all normal worlds but are highly probable are epistemically on par.

If we accept Parity, then the following seems to be a natural way of dealing with cases in which the two ways of being justified point in different directions: if P is true in all normal worlds and \sim P fails to be true in all normal worlds but is highly probable, then neither P nor \sim P are justified. Here is a reason for thinking that this may be the case. The fact that P is true in all normal worlds undermines justification for \sim P, despite the fact that \sim P is highly probable, and vice versa; the fact that \sim P, while not being true in all normal worlds, is highly probable undermines the justificatory status of P, despite the fact that P is true in all normal worlds. Following this strategy, in cases where the two ways of being justified are in conflict, the agent should suspend judgment. This idea can be captured in the following norm.

Suspension If, given S's evidence, the belief that P is true in all normal worlds, and the belief that \sim P fails to be true in all normal worlds but is highly probable for S, then S lacks justification for both P and \sim P and should suspend judgment.

Accepting Suspension leads to the following solution to the preface paradox. The author would lack justification for both (P2), the claim that her book is error-free, and (P3), the claim that the book contains at least one error. Instead, Suspension suggests that the author should suspend judgment about whether her book contains any errors. Since this solution would reject both (P2) and (P3) it avoids the contradictory conclusion in (P4). This solves the paradox.

While the above strategy assigned equal weight to the different ways of being justified this is by no means the only option. Another option for resolving tensions between different ways of being justified is to assign different weight to different justification-realizing properties. For instance, one might want to assign more weight to beliefs formed in a Risk Accumulating Epistemic Environment and that are highly probable than to beliefs formed in a Normal Epistemic Environment. In this case we would treat beliefs that are not true in all normal worlds but that are nevertheless highly probable as epistemically privileged over beliefs that are true in all normal worlds. Let's call this proposal *Priority-P*.

Priority-P Beliefs that fail to be true in all normal worlds but are highly probable are epistemically privileged over beliefs that are true in all normal worlds.

Priority-P might then be used to underwrite the argument that in cases where beliefs that fail to be true in all normal worlds but are highly probable and beliefs that are true in all normal worlds are at odds – i.e. where two justification-conferring properties support contradictory beliefs - preferential treatment should be given to beliefs that are highly probable. Following this proposal, one might insist that if with regards to P one is in a Risk Accumulating Epistemic Environment and P is highly probable and with regard to $\sim P$ one is in a Normal Epistemic Environment such that $\sim P$ is true in all normal worlds, then the high probability of P defeats justification for $\sim P$ (but not *vice versa*). Subsequently, S would be justified in believing P, the proposition that fails to be true in all normal worlds but is highly probable, and lack justification for believing $\sim P$, the proposition that is true in all normal worlds. Thus Priority-P might be used to motivate a defeat norm along the following lines.

Defeat-P If P fails to be true in all normal worlds but is highly probable for S, and $\sim P$, given the evidence, is true in all normal worlds, then S is justified in believing P and lacks justification for believing $\sim P$

With this norm in place the pluralist account of justification would secure Christensen's (and other opponents' of (MPC)) preferred solution to the preface paradox, namely that the author lacks justification for believing that the book is error-free. Why? Recall that, on the evidence, the belief that the book is error-free is true in all normal worlds. In contrast, the belief that the book contains at least one error, while failing to be true in all normal worlds, is highly probable. Since according to Defeat-P beliefs that fail to be true in all normal worlds but are highly probable undermine the justificatory status of beliefs that are true in all normal worlds, the belief that there is at least one error in the book will be justified, and, moreover, it will defeat the author's justification for believing that the book is error free. So, the preface paradox is solved by denying (P2).

A final option would be to reverse the priority ranking suggested above such that beliefs formed in a Normal Epistemic Environment and that are true in all normal worlds would be epistemically privileged over highly probable beliefs formed in Risk Accumulating Epistemic Environments. This way of weighing the two justification-conferring properties can be captured as follows.

Priority-N Beliefs that are, on the evidence, true in all normal worlds are epistemically privileged over beliefs that fail to be true in all normal worlds but are highly probable for S.

If we accept Priority-N, then we might say the following about cases where the two justification-conferring properties would lead to contradictions. In cases where beliefs that are true in all normal worlds and beliefs that fail to be true in all normal worlds but that are highly probable are in conflict, priority should be given to beliefs that are true in all normal worlds. As such, in cases where a belief that P is true in all normal worlds and the belief that $\sim P$ fails to be true in all normal worlds but is highly probable for S, the fact that P is true in all normal worlds undermines S's justification for believing $\sim P$. In these cases then, S would be justified in believing P and would lack justification for believing $\sim P$. This approach, on which the justification-conferring property of *being normally true* is taken to be epistemically privileged over the justification-conferring property of *being highly probable* would support a defeat norm along the following lines.

Defeat-N If the evidence normically supports P and $\sim P$ fails to be true in all normal worlds but is highly probable for S, then S is justified in believing P and lacks justification for believing $\sim P$

This proposal would deal with the preface paradox as follows. According to our analysis above, the belief that the book is error-free is true in all normal worlds. In contrast, the belief that there is at least one error in the book fails to be true in all normal worlds but is highly probable. According to Defeat-N, the fact that the belief that the book is error-free is true in all normal worlds defeats the author's justification

for believing that there is at least one error in the book, despite the fact that this is highly probable. As a result, the author is justified in believing that the book is error-free and lacks justification for believing that the book contains at least one error. This solution would solve the preface paradox by denying (P3) and thereby secure the result favoured by Pollock (1986), Ryan (1991), Kaplan (1996, 2013), Leplin (2009) and Smith (2010, 2016) and other proponents of (MPC).

Which of these three options should we endorse? It is important to point out that the functional theory is compatible with all three solutions and by itself does not commit us to any particular one. If one wanted to, one could simply endorse the option that would deliver one's preferred solution to the preface paradox. There is however a worry that this is not a particularly principled way of deciding this issue. Just imagine the following exchange, "Why do you think we should endorse Priority-N over Priority-P?" "Oh, it just secures my preferred solution to the preface paradox." Even though I do not want to rest too much on this issue, I take it that it is rather uncontroversial that there is something unsatisfactory about this response. Surely it would be preferable if one had an *independently motivated* reason for accepting one of the strategies. Now, while I continue to remain officially neutral on which strategy we should endorse, I think that there is a persuasive and independently motivated reason pointing in favour of the last solution - i.e. the solution that assigns priority to beliefs that are true in all normal worlds and accepts *Priority-N* and *Defeat-N*.

4.3 A Quick Argument in Favour of Priority-N and Defeat-N

Note that in earlier discussion (chapter 2), we observed that when we compare beliefs based on, for instance, perception, testimony, or memory, with beliefs based on purely statistical evidence, then intuitively - regardless of how probable the statistical evidence makes P - many consider the former beliefs epistemically superior to the latter. Now, recall that for beliefs based on perception, testimony, or memory, we are usually in a Normal Epistemic Environment and the property that makes these beliefs justified is the property of *being normally true*. In contrast, for beliefs based on statistical evidence we are in a Risk Accumulating Epistemic Environment and the property that makes these beliefs justified is the property of *being highly probable*.

The fact that beliefs based on perception, testimony, and memory strike us as epistemically superior to beliefs based on purely statistical evidence, points towards the following more general claim: intuitively we consider beliefs that have the property of being true in all normal worlds to be epistemically privileged over beliefs that lack this property even if they are highly probable.

To further support this more general claim we may consider additional cases of justified beliefs formed in Risk Accumulating Epistemic Environments in which the evidence is not purely statistical like Helen's beliefs in *Lightning* and *Allergy*. It seems reasonable to think that we would still judge beliefs based on perception, testimony, and memory to be epistemically superior compared to Helen's beliefs in *Allergy* (that the brownie is safe to eat it) or *Lightning* (that she will see Bob next weekend). These considerations seem to lend further support to the claim that ordinarily we regard justified beliefs formed in Normal Epistemic Environments as being epistemically superior to justified beliefs formed in Risk Accumulating Epistemic Environments. So, if we want to take seriously the intuitive data about the comparative strength between beliefs that are true in all normal worlds and beliefs that do not have this property but are nevertheless highly probable, then we have reason to accept *Priority-N* and *Defeat-N*.

The solution to the preface paradox that arises from *Priority-N* and *Defeat-N*, as argued previously, is that author is justified in believing that the book is error-free and lacks justification for believing that it contains at least one error. Importantly however, following the argument outlined above, this result is now underwritten by more general observations about our ordinary judgments about the comparative strength between beliefs that are true in all normal worlds and belief that, despite being highly probable, will be false in some normal worlds. In short, this solution to the preface paradox now has some independent motivation.

Now, I am not claiming that this argument is conclusive, and, to reiterate, I remain officially neutral on the question of which of the three solutions to the preface paradox we should accept. However, it is important to point out that the preface paradox, regardless of whether we are pluralists or monists about justification, will inevitably require us to make a choice – it is a paradox after all. Just like the monist, who needs to decide whether or not they find the inductive reasons in support of the

claim that the book contains at least one error sufficient for justification, so the pluralist has to decide what to do about cases in which the two different ways of being justified are in conflict and whether or not one of the two justification-conferring properties is privileged over the other. And for both, the pluralist as well as the monist, regardless of what position one takes, there is bound to be some controversy surrounding it. Hence, the fact that an acceptance of *Priority-N* and *Defeat-N* is going to be controversial should not surprise us. But the fact that it is underwritten by more general reflections on how we ordinarily think about comparisons between beliefs formed in Normal Epistemic Environments and beliefs formed in Risk Accumulating Epistemic Environments, should make it *prima facie* more attractive than its alternatives. At the very least, it should make it reasonable to stand by this solution until a more convincing defence of one of the alternative proposals (or an altogether new proposal) has been produced.

4.4 Why is the Preface Paradox So Difficult?

A satisfactory solution to a paradox does not just require a story about how to respond to the paradox, it also requires a story as to what got us into the paradox in the first place and, in the case of a particularly intractable paradox, a story as to why the paradox was so intractable. In what follows I argue that the functional theory does not just offer a very compelling explanation as to why the preface paradox is so resilient, but also that this explanation is more convincing than the one offered by monists.

In explaining why we feel so inclined to accept both (P2) – the claim that the author is justified in believing that her book is error free - and (P3) – the claim that book contains at least one error - the monist will typically tell us that we were simply *mistaken* about what justification requires. For instance, Kaplan and Kim will tell us that we felt inclined to accept (P3) because we *mistakenly* thought that a belief could be justified in the absence of a *special reason*. Similarly, proponents of normalcy views like Smith and Leplin will tell us that we felt inclined to accept (P3) because we *mistakenly* thought that high probability is sufficient for justification. Alternatively, proponents of the Lockean View will tell us that we felt inclined to accept (P2) because we *mistakenly* thought that (MPC) was a valid principle. In short,

monists will usually tell us that what gets us into the preface paradox is that we are somehow widely mistaken about some feature of the notion of justification.

But is this explanation compelling? Here is a reason to think that it is not. Normally, when it is pointed out to us that we are mistaken about something, we simply correct our beliefs and move on. But this means that in the case of the preface paradox we should expect that after it is pointed out to us that we were mistaken about one of the premises, the paradox should lose its force. In other words, we should no longer feel inclined to accept both (P2) and (P3); after all, we just learned that we were mistaken about at least one of them. That this actually happens however is questionable. For many the preface paradox remains puzzling even after learning about our alleged mistake, and it is this continued pull towards both (P2) and (P3) that is difficult to explain if it really was the case, as the monist suggest, that we were simply mistaken about one of them. It seems then that the monist has some difficulties in explaining why the preface paradox is so intractable.

The functional theory, I argue, offers a more nuanced and ultimately more compelling explanation of the intractability of the preface paradox. Importantly, this explanation does not require that we were widely mistaken about justification.

Recall our earlier diagnosis of what's going on in the preface paradox. In the preface paradox the two ways of being justified pull in different directions and make contradictory predictions. While this diagnosis is straightforward, it also points towards the following important thought that may easily be missed: if it wasn't for the fact that the two beliefs (that the book is error-free and that it contains at least one error) give rise to a contradiction, both beliefs *would be* justified according to the functional theory. I find this thought to be very persuasive. But neither the Lockean View nor the Normalcy View can accommodate this idea. On the functional theory then, it is not that we've all along been *mistaken* about what justification requires, but rather that the two incompatible beliefs in (P2) and (P3) are both *prima facie* justified. This provides a compelling explanation for the intractability and stubbornness of the preface paradox.

5. *Legal Cases*

Finally, we will consider how the functional theory will deal with the puzzles from legal cases considered earlier in chapter 2. Let's first consider the standard case we previously called Prisoners before looking at the more difficult comparative cases – cases in which we compare the epistemic status of beliefs based on purely statistical evidence with beliefs formed using other types of evidence such as perception or testimony. In chapter 2 we presented Prisoners as follows.

Prisoners 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The 100th prisoner played no role in the assault and could have done nothing to stop it. There is no further information that we can use to settle the question of any particular prisoner's involvement.

After the attack, the prison officials decide to legally pursue a randomly chosen prisoner from the yard for their alleged involvement in the attack. This case (and other similar ones) give rise to at least two interesting questions: first, is it permissible to convict the defendant? Here people almost universally agree that the answer is, no. Following Littlejohn (forthcoming) we said that many, including legal scholars, will want to reject Punish.

Punish It is permissible to punish the defendant in Prisoner (and similar cases where the only evidence of guilt is statistical evidence).

And second, is it permissible to believe that the defendant is guilty? Note that the answer to this epistemic question is less obvious than the answer to the previous question. Assuming for the sake of argument that it is, Littlejohn called this Believe.

Believe It is permissible to believe that the defendant is guilty in Prisoners (and similar cases where the only evidence of guilt is statistical evidence).

Now what do different monist theories say about the status of Believe? As argued in chapter 2, since the proposition that the prisoner was involved in the attack is highly

probable, the Lockean View will predict that we are justified in believing it. Hence, proponents of the Lockean View are likely to accept Believe.

Things are different for proponents of the Normalcy View (and the Strong View). Since on the evidence there exists at least one normal world in which we falsely believe that the prisoner was involved in the attack – this is the world in which the randomly chosen prisoner just happened to be the one prisoner in the yard who was not involved in the attack - the belief would not be justified. As a result, proponents of the Normalcy View will reject Believe.

What does the functional theory predict? On the functional theory we will be justified in believing that the prisoner was involved in the attack. Why? First consider in what epistemic environment we are with regards to the proposition P, that the prisoner is guilty. Since there exists at least one normal world in which we falsely believe that P, we are in a Risk Accumulating Epistemic Environment with regards to P. This means that the belief will be justified as long as it is highly probable on the evidence. Since this is the case for P, the belief that the defendant was involved in the attack will be justified on the functional theory. Thus, like proponents of the Lockean View, proponents of the functional theory will accept Believe.

Now, does the acceptance of Believe also commit us to accept Punish? If this were to be the case, then, since Punish seems unacceptable, the functional theory would be in trouble. Fortunately, however, as argued in chapter 2, accepting Believe does not commit us to accept Punish. This, we argued, is the case because the two questions, whether it is permissible to believe someone to be guilty of some crime and whether it is permissible to find them guilty in court, reasonably come apart. While the first question is primarily an epistemic one, the second question has a distinctly moral dimension. Hence, we argued that one might accept Believe but nonetheless reject Punish on moral grounds. This makes room for the following position: while it is epistemically permissible for us to believe that the evidence is guilty, it is nevertheless impermissible to convict the defendant because doing so would violate one or more of our moral ideals concerning the just and fair treatment of defendants. Hence, the functional theory does not require that we accept Punish and subsequently Prisoners does not present a problem for the functional theory. Next, we will take a look at how the functional theory can deal with the more difficult comparative cases.

5.1 Comparative Cases

In chapter 2 we also introduced the following variant of Prisoners.

Prisoners* 100 prisoners are exercising in the prison yard. Suddenly a prisoner attacks one of the guards. After the attack, one of the guards claims to have recognized the responsible prisoner – let’s call him Bill – and offers to testify as an eyewitness in court. The environmental conditions on the day of the attack were just as they are on most other days.

In Prisoners* the correct answers to our two initial questions - whether it is permissible to convict Bill and whether it is permissible to believe that Bill was responsible for the attack - are rather uncontroversial. In response to the first question, most people would judge that in light of the eyewitness testimony it is permissible to find Bill guilty. Things are similar for the second question. According to most prominent theories of epistemic justification we will be justified in believing that Bill is guilty. Let’s briefly consider why.

Let’s follow our previous convention from chapter two and call the proposition that the randomly chosen defendant in Prisoners was involved in the attack P, and the proposition that Bill in Prisoners* attacked the guard P*. On the Lockean View, the story is straightforward. According to the Lockean View we are justified in believing P* because in light of the eyewitness testimony P* is highly probable. What about the Normalcy View? According to the Normalcy View we will also be justified in believing P*. Why? As in the case of lotteries, it seems plausible that given the evidence there does exist a normal world in which we falsely believe P*. To illustrate this idea, we may consider Smith’s notion of calling for an explanation: if it turned out that despite the guard’s eyewitness testimony our belief that P* turned out to be false, then this would call for some special explanation. Some mitigating or interfering circumstances would need to obtain to explain why we falsely believe that Bill was responsible for the attack – maybe the guard was tired and mistook Bill for someone else, perhaps there was another prisoner which looked just like Bill from the angle at

which the guard was overlooking the prison yard, or maybe the guard was plotting against Bill.

What about The Functional Theory? First, let's consider in what epistemic environment we are with regards to P*. Since there does not exist a normal world in which we falsely believe that P*, we are in a Normal Epistemic Environment. And since the belief that P* has the property of *being normally true*, it will be justified on the functional theory.

However, the job is not done here. For those views that have predicted that in both Prisoners and Prisoners* we are justified in believing that the defendant was involved in the attack – i.e. for the Lockean View and the Functional Theory - there remains a challenge. Many have the intuition that in Prisoners* our epistemic status with regards to P* is somehow privileged compared to P in Prisoners; or, that our justification for P* is somehow stronger than our justification for P. In chapter two we called this the Comparative Claim.

The Comparative Claim We are more justified in believing P* than P.

An important question then is whether the Functional Theory of Justification, which is our focus here, can somehow accommodate The Comparative Claim. I will turn to this now.

Recall that for the Lockean View this was a difficult challenge. As argued in chapter 2, Lockeans will likely need to reject The Comparative Claim and explain away the intuitive pull to accept it. Let's briefly review why Lockeans will struggle to accommodate The Comparative Claim. First consider the degree of probability that we should assign to the proposition P in Prisoners. Given our evidence, or the belief-forming method employed, the degree of probability that we should assign to the proposition that the randomly chosen prisoner was involved in the attack, is 0.99; $\Pr(P)=.99$. Now what degree of probability should we assign to P* in Prisoners*? Now, regardless of what we think the precise numeric value should be, what's important is that it will be less than 1; after all, P* is not a necessary truth and could possibly be false. We previously considered some of the circumstances that would explain why P* may be false: maybe the guard was tired and mistook Bill for someone

else, perhaps there was another prisoner which looked just like Bill from the angle at which the guard was overlooking the prison yard, or maybe the guard was plotting against Bill. Let's assume, for the sake of argument, that degree of probability we should assign to P^* is 0.98; $\Pr(P^*)=0.98$.¹³⁹ Here then is the puzzle for proponents of the Lockean View: since $\Pr(P)=0.99 > \Pr(P^*)=0.98$, proponents of the Lockean View seem committed to the idea that we are more justified in believing that the randomly chosen defendant in Prisoners was involved in the attack than we are in believing that Bill attacked the guard in Prisoners*, despite the fact that we have eyewitness testimony against Bill.¹⁴⁰ But this runs against our intuitions.

How can Lockeans respond to the challenge that they cannot accommodate the intuitions underwriting The Comparative Claim? As we have seen in chapter 2, the most promising strategy for Lockeans is to reject Comparative Claim and provide an explanation for why we should reject our intuitive judgment that P^* is more justified than P . One plausible option we considered is the following. Since a commitment to the Lockean View often comes with a commitment to the idea that justified belief aims at truth (see ch. 2), they may argue that according to The Comparative Claim we should regard a proposition that is less probable (i.e. P^*) to be more justified than a proposition that is more likely to be true (i.e. P). And, if we accept that justified belief aims at truth, then we have good reason to reject this claim. While this may provide Lockeans with a perhaps reasonable explanation for why we should reject The Comparative Claim, we may nevertheless find this consequence of their view unpersuasive. After all, the intuition that our belief in Prisoners* is more justified than our belief in Prisoners is very strong. In chapter 2 we called this challenge for the Lockean View the *Residual Problem*.

¹³⁹ The precise degrees of probability here are irrelevant; As long as $\Pr(P^*) < 1$ we can adjust the number of prisoners to match (or exceed) the degree of probability one thinks we should assign to P^* . Hence as long as $\Pr(P^*) < 1$ we will be able to get the puzzle off the ground.

¹⁴⁰ In chapter two we argued that to run this argument, we strictly speaking require a further principle along the following lines.

Probabilistic Comparison For any two justified beliefs P and Q , if P 's degree of probability exceeds Q 's degree of probability, then P is more justified than Q .

For our discussion here, I will simply assume that some such principle will hold.

Let's now turn our attention to the functional theory and see whether it is able to somehow accommodate The Comparative Claim. First, it is important to consider what's involved in the comparison between P and P*. On the functional theory, the two propositions P and P*, while both justified, are justified in different ways. With regards to P we are in a Risk Accumulating Environment, such that the proposition that the randomly chosen prisoner was involved in the attack is justified in virtue of *being highly probable*. In contrast, with regards to P* we are in a Normal Epistemic Environment such that the belief that Bill attacked the guard is justified in virtue of *being normally true*. Now, as in the preface paradox, the question of whether one of these two beliefs is more justified than the other reduces to the more general question of what we want to say about the comparative status of two beliefs that are justified in different ways. As we have seen in the discussion of the preface paradox, there are a number of options.

If we wanted to follow the Lockean View and deny The Comparative Claim – i.e. deny that P* is more justified than P – then this result can be achieved by opting for what we called Parity in our earlier discussion of the preface paradox. According to Parity, the different ways of being justified are epistemically on par. If we accept this claim, then it is natural to think that when comparing two beliefs that are justified in different ways, neither of them would be more justified than the other. In other words, from Parity there is a natural way to the following comparison norm.

Comparison If S's belief that P is true in all normal worlds and S's belief that Q fails to be true in all normal worlds but is highly probable for S, then P and Q are equally justified for S.

In the case of Prisoners and Prisoners*, this would mean that we have equal justification for both P and P* and neither belief would be more justified than the other. So, if we accept Parity and Comparison, then the functional theory will not accommodate The Comparative Claim. This however is not the only option available to us.

Let's suppose that unlike the Lockean we want to take seriously our intuitive judgments about the comparison between P and P* and accommodate The

Comparative Claim. On the functional theory, what's underlying The Comparative Claim is the more general idea that justified beliefs formed in a Normal Epistemic Environment – i.e. beliefs that, on the evidence, are true in all normal worlds – are somehow epistemically privileged, or superior, to justified beliefs formed in Risk Accumulating Environments – i.e. beliefs that fail to be true in all normal worlds but are highly probable for the subject. We previously captured this idea in the principle *Priority-N*. And from *Priority-N* there is natural way to the following comparison norm, according to which beliefs that are true in all normal worlds are more justified than beliefs that fail to be true in all normal worlds even if they are highly probable.

Comparison-N If S's belief that P is true in all normal worlds and S's belief that Q fails to be true in all normal worlds but is highly probable for S, then S's belief that P is more justified than S's belief that Q.

Comparison-N would predict that our belief that Bill in Prisoners* was responsible for the attack is more justified than our belief that the randomly chosen prisoner in Prisoners was involved. This is the intuitively correct result. So, if we think that a theory of justification should capture the intuitive data expressed in The Comparative Claim, then the functional theory can deliver this result via a comparison norm like Comparison-N.¹⁴¹

However, it is important to note that *Priority-N* and *Comparison-N* are only *optional* features of the functional theory and that we can therefore stay officially neutral on their status. But, as we have seen, on the functional theory there are a number of different options for dealing with the more general problem of how to

¹⁴¹ There is of course a third possible option. Following *Priority-P*, which assigns more weight to beliefs formed in a Risk Accumulating Epistemic Environment and that are justified in virtue of being highly probable, we could motivate a comparison norm according to which P comes out as more justified than P*.

Comparison-P If S's belief that P fails to be true in all normal worlds but is highly probable for S and S's belief that Q is true in all normal worlds, then S's belief that P is more justified than S's belief that Q.

However, accepting *Comparison-P* would not just fail to accommodate The Comparison Claim, it would also yield the result that P in Prisoners is more justified than P* in Prisoners*. This is effectively the result that caused some trouble for the Lockean View. Since it is difficult to see the motivation for this proposal, I will not consider it further here.

compare two (or more) beliefs that are justified in different ways; and the functional theory does not commit us to any particular candidate.

5.2 A Brief Comparison

Finally, let's briefly compare the three theories – the probabilistic Lockean View, the non-probabilistic Normalcy View, and the pluralist Functional Theory of Justification - with regards to their ability to accommodate The Comparative Claim. In the end, I think there are good reasons to think that the functional theory occupies the most attractive position regarding the comparison between P in Prisoners and P* in Prisoners*.

On the Lockean View, we were unable to accommodate The Comparative Claim. Since P* in Prisoners*, if true, will only be so contingently, its degree of probability will be less than 1, and since in Prisoners we can always adjust the numbers of prisoners in the yard, we will always be able to set up a case in which $\Pr(P) \geq \Pr(P^*)$. As a result, we will always be able to generate a case that violates The Comparative Claim. As we have seen, the perhaps most compelling way for the Lockean to deal with this problem will be to try and explain why denying The Comparative Claim, while counterintuitive, is actually the correct result.

On the Normalcy View The Comparative Claim is trivially accommodated. P* is more justified than P because P was never justified to begin with. However, here is a worry about this result. The intuitive data that we set out to capture was that P* is *more justified* than P – i.e. that we are *more justified* in believing that Bill in Prisoners* was responsible for the attack than we are in believing that a randomly chosen prisoner in Prisoners was involved in the attack. But it is unclear that the intuitive data also supports the stronger claim that P is altogether unjustified. In fact, the empirical work on beliefs based on statistical evidence – usually lottery beliefs - suggests that people do tend to judge that beliefs based on strong statistical evidence can be justified.¹⁴² Hence, we might worry that accommodating our intuition that P* is *more justified* than P by denying that P is justified in the first place is too strong. In other words, it

¹⁴² See, Turri and Friedman (2014), Friedman and Turri (2015), Ebert et al. (2018).

seems like the Normalcy View, in order to accommodate The Comparative Claim, may overshoot the target.

On the Functional Theory, unlike on the Normalcy View, both P and P* will be justified. So, we can avoid the worry of being able to accommodate The Comparative Principle only at the cost of making potentially counterintuitive claims about the justificatory status of P. And unlike the Lockean View, the functional theory is able to accommodate The Comparative Claim if we want to. Since on the functional theory we need to decide how to compare beliefs that are justified in different ways, we can accommodate The Comparative Claim, if we think that beliefs that are true in all normal worlds are epistemically privileged over beliefs that fail to be true in all normal worlds but are nevertheless highly probable. As such, our judgment underwriting The Comparative Claim could be used to motivate principles like *Priority-N* and *Comparison-N*.

Thus, the functional theory seems to provide a more compelling way of dealing with comparative cases like the one involving Prisoners and Prisoners*. Unlike the Lockean View it is able to secure the intuitively correct result that P* is more justified than P without having to follow the path of the Normalcy View which can only secure this intuitively correct result at the cost of making another, perhaps counterintuitive, prediction about the justificatory status of one the beliefs in question - i.e. P.

6. *Taking Stock*

In the previous chapter we introduced a pluralist theory of justification - the Functionalist Theory of Justification – as a way of avoiding the (MPC)-Stinginess Dilemma faced by its monist counterparts. In this chapter I have argued that the functional theory also provides interesting new solutions to a number of epistemic puzzles and paradoxes. As we have seen, the new resources offered by a theory on which a belief can be justified in more than one way (or by more than one property) allow us to avoid many of the problems faced by theories that insist that there is only one property that can justify belief.

One broad family of epistemic paradoxes aim to show that justification is not closed under multi premise deduction. Members of this family include the Lottery Paradox, the Preface Paradox, and the new Paradox of the Pill. In this chapter we have seen that these paradoxes rely in an important way on the assumption that there is one and only one way for a belief to be justified; for as soon as we part with this assumption these paradoxes lose much of their initial force. The most stubborn of these paradoxes is perhaps the preface paradox. However, we have seen that the functional theory is not just compatible with a number of different solutions to the preface paradox but that it also provides an interesting and compelling diagnosis of both the nature of the paradox and its intractability.

Similarly, the functional theory provides compelling ways of dealing with the epistemic puzzles surrounding legal cases involving statistical evidence. Especially when it comes to the comparative cases and the accommodation of The Comparative Claim the functional theory appears to perform better than its monist competitors. Where the Lockean View is unable to accommodate The Comparative Claim and the Normalcy View can only accommodate it at the expense of denying that P in Prisoners is justified, the functional theory puts us in a position to accommodate The Comparative Claim without having to deny the justificatory status of P.

In conclusion, the pluralist Functional Theory of Justification does not just help us avoid the uncomfortable (MPC)-Stinginess Dilemma, arguable one of the most pressing obstacles we face in providing a satisfactory theory of justification, but it also has the consequence of making many epistemic puzzles and paradoxes appear less daunting.

Conclusion

So, what have we done? We started out by considering the general methodological question of how to determine whether some theory is a theory of justification rather than something else. Put in terms of properties, we set out to answer the question of how we may go about determining whether some property is a property that makes a belief justified – i.e. whether it is a *justification-conferring* property or a *J-property*, as we called them - rather than something else? In response to this question we proposed what I dubbed the *Minimalist Concept of Justification*, according to which a property is a candidate J-property iff it satisfies justification's core platitudes. The three core platitudes we identified were *Truth Candidacy*, *Permission*, and *Blamelessness*.

We then looked at two very different proposals about what this J-property may be. First, we considered the *Lockean View*; the perhaps most widely endorsed picture of justification in contemporary epistemology. According to the Lockean View the property that makes a belief justified is the property of *being highly probable*. The second broad family of views we considered was what I called the *Normalcy View*; a non-probabilistic alternative to the Lockean View. According to the Normalcy View, the property that makes a belief justified is the property of *being true in all normal worlds*; or, for short, of *being normally true*. However, ultimately, we found both proposals wanting.

Probabilistic accounts following the Lockean View are incompatible with the attractive principle of *multi premise closure* (MPC), while accounts following the Normalcy View, due to their modal strength, will be too demanding and therefore too stingy. This we argued leaves us in a dilemma - we called it the *(MPC)-Stinginess Dilemma*. The *(MPC)-Stinginess Dilemma* should worry epistemologists. If we think that the probabilistic/non-probabilistic divide presents a fundamental choice point in the justification debate - and as far as I can see there is no reason to deny that it is - then the fact that this divide leaves us in a dilemma is surely concerning. In fact, I am

inclined to think that the (MPC)-Stinginess Dilemma is one of the most pressing obstacles we face in providing a satisfactory theory of justification.

The question that naturally arose at this point was how we should respond to this dilemma. Finding any of the immediately available solution unsatisfactory, I proposed that we abandon J-monism – the idea that there is only one property that can make a belief justified. Instead I argued we should endorse J-pluralism – the idea that there is more than one property that can make a belief justified. With J-pluralism in place we were now able to say that some beliefs are justified in virtue of *being normally true*, while other beliefs are justified in virtue of *being highly probable*. An important upshot this picture is that for all beliefs that are justified in virtue of being normally true (MPC) will hold, while beliefs that lack this property may still be justified as long as they are highly probable. This allowed us to sidestep the dilemma.

Now, to turn pluralism about justification into a workable theory of justification, we needed to provide an overarching metaphysics that tells us, amongst other things, how the different notions of justification interact. For this we drew inspiration from previous work by *alethic pluralists* like Crispin Wright and Michael Lynch. Following Lynch's recent proposal of a functional theory of truth I proposed and defended what we called *The Functional Theory of Justification*. According to the functional theory, the property of *being justified* is a functional property, which can be realized (or manifested) by more than one property depending on our *epistemic environment*. The metaphysical backbone of the theory was provided by the *Minimalist Conception of Justification* introduced earlier. As it turns out there is simply more than one property that satisfies justification's core platitudes.

Finally, we argued that aside from helping us avoid the (MPC)-Stinginess Dilemma, the functional theory also provides attractive and interesting new solutions to a number of epistemic puzzles and paradoxes; these include the Lottery Paradox, the Preface Paradox, the newly proposed Paradox of the Pill, as well as a number of justification puzzles involving statistical evidence and legal cases. As we have seen, with two justification-conferring properties working side-by-side, the functional theory was able to offer more nuanced and ultimately more satisfactory solutions to these puzzles compared to its monist competitors.

Overall then, the *Functionalist Theory of Justification* has been shown to provide an attractive and promising new framework for thinking about epistemic justification. The strength of the theory ultimately derives from its effort to take seriously the epistemic complexity and diversity of the world in which we live.

Bibliography

- Alston, W. P. (1985). Concepts of Epistemic Justification. *The Monist* 68 (2):57-89.
- Alston, W. P. (1988). An Internalist Externalism. *Synthese* 74 (3): 265 - 283.
- Anderson, C. (2017). Between Probability and Certainty: What Justifies Belief By Martin Smith. *Analysis* 77 (3): 670–672.
- Arpaly, N. (2002). *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.
- Backes, M. (forthcoming). A Bitter Pill for Closure. *Synthese*.
- Backes, M. (forthcoming). Normalcy, Justification, and the Easy-Defeat Problem, *Philosophical Studies*.
- Baumann, P. (2014). Justification and the Truth-Connection By Clayton Littlejohn. *Analysis* 74 (4):731-733.
- Beall, J. C. (2000). On mixed inferences and pluralism about truth predicates. *Philosophical Quarterly* 50 (200):380-382.
- Bergmann, Michael (2006). *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford: OUP.
- Bird, Alexander (2007). Justified judging. *Philosophy and Phenomenological Research* 74 (1):81–110.
- Blome-Tillmann, M. (2015). Sensitivity, Causality, and Statistical Evidence in Courts of Law. *Thought: A Journal of Philosophy* 4 (2):102-112.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Bonjour, L. (2002). Intemalism and extemalism. In Paul K. Moser (ed.), *The Oxford Handbook of Epistemology*. Oxford University Press:234-263.
- BonJour, L. and Sosa, E. (2003). *Epistemic Justification: Internalism Vs. Externalism, Foundations Vs. Virtues*. Wiley-Blackwell.
- Brown, J. (2013). Impurism, Practical Reasoning, and the Threshold Problem. *Noûs* 47(1): 179-192.

- Brown, J. (2017). Blame and wrongdoing. *Episteme* 14 (3):275-296.
- Buchak, L. (2014). Belief, credence, and norms. *Philosophical Studies* 169(2):1-27.
- Chalmers, D. (2011). Verbal Disputes. *Philosophical Review* 120 (4):515-566.
- Chan, T. (2013). Introduction: Aiming at Truth. In *The Aim of Belief*. Oxford: OUP.
- Christensen, D. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: OUP
- Christensen, D. (2007). Three Questions about Leplin's Reliabilism. *Philosophical Studies* 134 (1): 43 - 50.
- Chuard, P. & Southwood, N. (2009). Epistemic norms without voluntary control. *Noûs* 43 (4):599-632.
- Coffman, E. (2010). Is Justified Belief Knowledge? Critical Notice of Jonathan Sutton, Without Justification. *Philosophical Books* 51:1-21
- Cohen, L. J. (1977). The probable and the provable. Oxford: OUP.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies* 46 (3):279-295.
- Colyvan, M., Ferson S., and Regan, H. (2001). Is it a crime to belong to a reference class? *The Journal of Political Philosophy* 9(2): 168-181.
- Comesaña, J. (2009). What lottery problem for reliabilism? *Pacific Philosophical Quarterly* 90(1): 1-20.
- Conee, E. & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies* 89 (1):1-29.
- Conee, E. & Feldman, R. (2001). Internalism defended. *American Philosophical Quarterly* 38 (1):1-18.
- Conee, E. and R. Feldman. (2004). *Evidentialism*. Oxford: OUP.
- Cruz, J. & Pollock, J. (2004). The Chimerical Appeal of Epistemic Externalism. In Richard Schantz (ed.), *The Externalist Challenge*. De Gruyter.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review* 111(2): 167-203.
- Douven, I. (2006). Assertion, knowledge, and rational credibility. *Philosophical Review* 115(4): 449-485.
- Douven, I. & Williamson, T. (2006). Generalizing the lottery paradox, *British Journal for the Philosophy of Science* 57 (4):755-779.

- Douven, I. (2008). The Lottery Paradox and our Epistemic Goal, *Pacific Philosophical Quarterly* 89 (2): 204-225.
- Easwaran K. and Fitelson B. (2015). Accuracy, Coherence, and Evidence, In T.S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology Vol. 5*. Oxford: OUP.
- Ebert, P., M. Smith, and I. Durbach. (2018). Lottery judgments: A philosophical and experimental study. *Philosophical Psychology* 31(1): 110-138.
- Eder, A.-M. A. (2015). No Match Point for the Permissibility Account. *Erkenntnis* 80 (3):657-673.
- Engel, P. (2005). Truth and the aim of belief. In D. Gillies (ed.), *Laws and Models in Science*. London: King's College Publications.
- Engel, P. (2013). In Defence of Normativism about the Aim of Belief. In T. Chan (ed.) *The Aim of Belief*. Oxford. OUP
- Enoch, D., T. Fisher, and L. Spectre. (2012). Statistical Evidence, Sensitivity, and the Legal Value of Knowledge. *Philosophy and Public Affairs* 40 (3):197-224.
- Enoch, D. and T. Fisher. (2015). Sense and sensitivity: epistemic and instrumental approaches to statistical evidence. *Stanford Law Review* 67: 557-611.
- Evnine, S. (1999) Believing Conjunctions, *Synthese* 118(2): 201-227.
- Fantl, J. and M. McGrath. (2002). Evidence, pragmatics, and justification. *Philosophical Review* 111(1): 67-94.
- Fantl, J. and M. McGrath. (2009). *Knowledge in an Uncertain World*. Oxford University Press.
- Feldman, R. (1988). Epistemic obligations. *Philosophical Perspectives* 2:235-256.
- Feldman, R. (2002). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Foley, R. (1979). Justified Inconsistent Beliefs, *American Philosophical Quarterly* 16(4): 247-257
- Foley, R. (1992). The Epistemology of Belief and the Epistemology of Degrees of Belief, *American Philosophical Quarterly* 29 (2): 111 - 124.
- Foley, R. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis, In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*. Springer.
- Friedman, O. and J. Turri. (2015). Is probabilistic evidence a source of knowledge? *Cognitive Science* 39: 1062–1080.

- Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues* 15(1):338–350.
- Goldman, A. (1979). What is Justified Belief ?, in *Justification and Knowledge*, ed. George Pappas, Boston, MA: D. Reidel: 1-25.
- Goldman, A. (1986). *Epistemology and Cognition*. Harvard University Press.
- Goldman, A. (1988). Strong and weak justification. *Philosophical Perspectives* (2):51-69.
- Goldman, A. (1999). Internalism exposed. *Journal of Philosophy* 96 (6):271-293.
- Graham, P. (2011). Does Justification Aim at Truth? *Canadian Journal of Philosophy* 41 (1):51-72.
- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
- Grundmann, T. (2015). Justification and the Truth-Connection, by Clayton Littlejohn. *Australasian Journal of Philosophy* 93 (3):622-624.
- Harman, G. (2002). Internal critique: A logic is not a theory of reasoning and a theory of reasoning is not a logic, In D. M. Gabbay et al. (eds.), *Handbook of the Logic of Argument and Inference The Turn Towards the Practical Vol. 1*. Amsterdam: Elsevier: 171 – 186.
- Hawthorne, J. (2003). *Knowledge and Lotteries*. Oxford: OUP.
- Hawthorne, J. & J. Stanley. (2008). Knowledge and Action. *Journal of Philosophy* 105(10): 571-590.
- Hirsch, E. (2005). Physical-object ontology, verbal disputes, and common sense. *Philosophy and Phenomenological Research* 70 (1):67–97.
- Horwich, P. (2013). Belief-Truth Norms. In T. Chan (ed.) *The Aim of Belief*. Oxford: OUP
- Ichikawa, J. J. (2014). Justification is Potential Knowledge, *Canadian Journal of Philosophy* 44 (2): 184-206.
- Jenkins, C. S. I. (2014). Merely Verbal Disputes. *Erkenntnis* 79 (1):11-30.
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge University Press.
- Kaplan, M. (2013). Coming to Terms with our Human Fallibility: Christensen on the Preface, *Philosophy and Phenomenological Research* 87 (1): 1-35.
- Kelp, C. (2011). Not Without Justification. *Dialectica* 65 (4):581-595.
- Kelp, C. (2017). Lotteries and Justification, *Synthese*:1-12

- Kim, B. (2015). This Paper Surely Contains Some Errors, *Philosophical Studies* 172 (4):1013-1029.
- Kroedel, T. (2012). The Lottery Paradox, Epistemic Justification and Permissibility, *Analysis* 72 (1): 57-60.
- Kvanvig, J. (2003). Propositional ism and the Perspectival Character of Justification. *American Philosophical Quarterly* 40 (1):3–18.
- Kvanvig, J. (2009). Assertion, knowledge, and lotteries. In Duncan Pritchard & Patrick Greenough (eds.), *Williamson on Knowledge*. Oxford: OUP.
- Kyburg, H. E. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press.
- Lackey, J. (2007). Norms of assertion. *Noûs* 41(4): 594–626.
- Latus, A. (2000). Our epistemic goal. *Australasian Journal of Philosophy* 78 (1):28-39.
- Lehrer, K. & Cohen, S. (1983). Justification, truth, and coherence. *Synthese* 55 (2):191-207.
- Lehrer, K. (1986). The coherence theory of knowledge. *Philosophical Topics* 14 (1):5-25.
- Leitgeb, H. (2017). *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.
- Leplin, J. (2009). *A Theory of Epistemic Justification*. Springer.
- Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, C. (forthcoming). Truth, Knowledge, and the Standard of Proof in Criminal Law. *Synthese*.
- Lynch, Michael P. (2004). Truth and multiple realizability. *Australasian Journal of Philosophy* 82 (3):384 – 408.
- Lynch, Michael P. (2009). *Truth as One and Many*. Clarendon Press.
- Madison, B. J. C. (2010). Epistemic Internalism. *Philosophy Compass* 5 (10):840-853.
- Makinson, D. C. (1965). The Paradox of the Preface, *Analysis* 25(6): 205-207.
- McGlynn, A. (2012). Justification as 'Would-Be' Knowledge. *Episteme* 9(4): 361-376.

- Moser, P. K. (1989). *Knowledge and Evidence*. New York, NY: Cambridge University Press.
- Nelkin, D. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review* 109 (3):373-409.
- Nelson, M. (2010). We Have No Positive Epistemic Duties. *Mind* 119 (473):83-102.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Papineau, D. (2013). There are no norms of Belief. In T. Chan (ed.) *The Aim of Belief*. Oxford. OUP
- Pedersen N. (2017). Pure Epistemic Pluralism. in Coliva A. & Nikolaj Jang Lee Linding Pedersen (eds) *Epistemic Pluralism*. Palgrave Innovations in Philosophy. Palgrave Macmillan.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: OUP.
- Pollock, J. L. (1983). Epistemology and Probability. *Synthese* 55 (2):231–252.
- Pollock, J. L. (1986). The paradox of the preface. *Philosophy of Science* 53 (2):246-258.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: OUP
- Pryor, J. (2004). What's wrong with Moore's argument? *Philosophical Issues* 14(1):349-378.
- Redmayne, M. (2008). Exploring the proof paradoxes. *Legal Theory* 14: 281-309.
- Reynolds, Steven L. (2013). Justification as the appearance of knowledge. *Philosophical Studies* 163(2): 367-383.
- Ross, J. and Schroeder, M. (2014). Belief, Credence, and Pragmatic Encroachment, *Philosophy and Phenomenological Research* 88 (2): 259-288.
- Ryan, S. (1996). The Epistemic Virtues of Consistency, *Synthese* 109(2):121-141.
- Ryan, S. (1991). The Preface Paradox, *Philosophical Studies* 64 (3):293-307.
- Shah, N. (2003). How truth governs belief. *Philosophical Review* 112(4):447-482.
- Shah, N. and Velleman, J. D. (2005). Doxastic Deliberation. *Philosophical Review* 114:497–534.
- Smith, M. (2010). What Else Justification Could Be, *Noûs* 44(1):10-31.
- Smith, M. (2014). Review of Justification and the Truth Connection by Clayton Littlejohn. *Philosophical Quarterly* 64:380-383.

- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: OUP.
- Smith, M. (forthcoming). When Does Evidence Suffice for Conviction? *Mind*.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* 13: 137-49.
- Staffel, J. (2016). Beliefs, buses and lotteries: Why rational belief can't be stably high credence. *Philosophical Studies* 173(7):1721-1734.
- Sturgeon, S. (2008). Reason and the Grain of Belief, *Noûs* 42(1): 139-165.
- Sutton, J. (2005). Stick to what you know. *Noûs* 39(3): 359–396.
- Sutton, J. (2007). *Without Justification*. MIT Press
- Swinburne, R. (2001). *Epistemic Justification*. Oxford: OUP.
- Tappolet, C. (1997). Mixed inferences: A problem for pluralism about truth predicates. *Analysis* 57 (3):209–210.
- Tappolet, C. (2000). Truth pluralism and many-valued logics: A reply to Beall. *Philosophical Quarterly* 50 (200):382-385.
- Thomson, J. J. (1986). Liability and individualized evidence. *Law and Contemporary Problems* 49(3): 199-219.
- Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and Phenomenological Research* 80 (2):312-326.
- Turri, J. & Friedman O. (2014). Winners and losers in the folk epistemology of lotteries. In J. Beebe (ed.), *Advances in Experimental Epistemology*. London: Bloomsbury.
- Vahid, H. (2010). Externalism/Internalism. In Sven Bernecker & Duncan Pritchard (eds.) *The Routledge Companion to Epistemology*. Routledge.
- Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford: OUP.
- Vogel, J. (1990). Are there Counterexamples to the Closure Principle?, in *Doubting: Contemporary Perspectives on Skepticism*, ed. Michael David Roth & Glenn Ross, Dordrecht: Kluwer Academic Publishers: 13 - 29.
- Vogel, J. (1999). The New Relevant Alternatives Theory, *Philosophical Perspectives* 13: 155-180.
- Wasserman, D. T. (1991). The Morality of Statistical Proof and the Risk of Mistaken Liability. *Cardozo Law Review* 13: 935 – 976.

- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives* 36(16): 267-97.
- Wedgwood, R. (2013). The Right Thing to Believe. In Timothy Chan (ed.), *The Aim of Belief*. Oxford: OUP.
- Weiner, M. (2005). Must we know what we say? *Philosophical Review* 114(2): 227-251.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives* 19 (1):445-459.
- Whiting, D. (2013). Nothing but the Truth: On the Norms and Aims of Belief. In Timothy Chan (ed.), *The Aim of Belief*. Oxford University Press.
- Williams, B. (1970). Deciding to believe. Reprinted in *Problems of the Self*. Cambridge: CUP.
- Williamson, T. (1994). A Critical Study of Crispin Wright, Truth and Objectivity. *International Journal of Philosophical Studies* 30 (1):130-144.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: OUP
- Wright, Crispin (1992). *Truth and Objectivity*. Harvard University Press.
- Wright, C. (1999). Truth: A Traditional Debate Reviewed. In Simon Blackburn & Keith Simmons (eds.), *Truth*. New York: OUP.
- Wright, C. (2013). A plurality of pluralisms. In Nikolaj Jang Lee Linding Pedersen & Cory D. Wright (eds), *Truth and Pluralism: Current Debates*. Oxford University Press.